arXiv:1610.05023v1 [physics.comp-ph] 17 Oct 2016

High-order absorbing boundary conditions with corner/edge compatibility for GPU-accelerated discontinuous Galerkin wave simulations

A. Modave*[1,2], A. Atle[3], J. Chan[1,4], and T. Warburton[1]

[1]Virginia Tech, Blacksburg, VA, USA
[2]POEMS (UMR 7231 CNRS-ENSTA-INRIA), ENSTA Paristech, Palaiseau, France
[3]TOTAL E&P, Houston, TX, USA
[4]Rice University, Houston, TX, USA

Abstract

Discontinuous Galerkin finite element schemes exhibit attractive features for accurate large-scale wave-propagation simulations on modern parallel architectures. For many applications, these schemes must be coupled with non-reflective boundary treatments to limit the size of the computational domain without losing accuracy or computational efficiency, which remains a challenging task. In this paper, we present a combination of high-order absorbing boundary conditions (HABCs) with a nodal discontinuous Galerkin method for cuboidal computational domains. Compatibility conditions are derived for HABCs intersecting at the edges and the corners of a cuboidal domain. We propose a GPU implementation of the computational procedure, which results in a multidimensional solver with equations to be solved on 0D, 1D, 2D and 3D spatial regions. Numerical results demonstrate both the accuracy and the computational efficiency of our approach. We have considered academic benchmarks, as well as a realistic benchmark based on the SEAM model used in exploration geophysics.

Keywords: wave propagation, unbounded domain, absorbing boundary condition, finite element, discontinuous Galerkin, time domain

1 Introduction

Numerical simulation tools play an important role for solving a wide range of large-scale wave-like problems in fields as diverse as underwater acoustics, electromagnetic scattering and seismic imaging. In this context, computational procedures based on discontinuous Galerkin finite element schemes are very attractive. These methods can provide accurate solutions to realistic wave-like problems thanks to heterogeneous, non-conforming and curvilinear meshes, high-order discontinuous basis functions, and stable formulations for complicated physical models (see *e.g.* [Cockburn et al., 2000; Hesthaven and Warburton, 2007; Kopriva et al., 2000; Léger et al., 2014; Li and Hesthaven, 2014; Mercerat and Glinsky, 2015; Schmitt et al., 2016; Toulorge and Desmet, 2010; Warburton, 2000, 2013; Wilcox et al., 2010; Ye et al., 2016]). In addition, the discrete structure of the numerical schemes is well suited for efficient massively parallel computing on distributed memory architectures and modern many-core accelerators (see *e.g.* [Chan et al., 2016b; Godel et al., 2010; Klöckner et al., 2009; Modave et al., 2015, 2016b; Seny et al., 2014]).

A critical issue for the simulation of wave phenomena is to correctly account for radiation of waves at artificial boundaries of the computational domain. Non-reflective boundary treatments must be incorporated into the discontinuous Galerkin formulations in order to simulate the outward propagation of signals and perturbations generated from within the computational domain, even if they are not *a priori* known. The

*Corresponding author: axel.modave@ensta-paristech.fr

challenge then consists of devising boundary treatments that preserve the accuracy of the numerical solution without overpenalizing the computational efficiency of the implementation.

Basic boundary techniques encompass characteristic-based conditions, impedance conditions and sponge layers. These techniques are robust, straightforward to implement and cheap to use, but they provide a relatively poor approximation of the solution. Two families of techniques provide high-fidelity solutions at reasonable computational cost: perfectly matched layers (PMLs) and local high-order absorbing boundary conditions (HABCs).

- With a PML, the computational domain is extended by a layer with an artificial medium that is both dissipative and perfectly matched with the original medium. Any outgoing wave is perfectly transmitted from the domain to the layer and damped by the medium. For transient problems, PML formulations are generally built by modifying the governing equations and introducing auxiliary fields inside the layer. The first PML was proposed by Bérenger [Bérenger, 1994] in computational electromagnetics for finite difference time domain simulations. Since this seminal work, a large number of PML formulations and discretizations have been proposed and analyzed for various wave problems (see *e.g.* [Appelö et al., 2006; Bécache et al., 2015; Bérenger, 1994; Bermúdez et al., 2010; Gedney, 1996; Hu, 2008; Komatitsch and Tromp, 2003; Lu et al., 2004; Modave et al., 2016a]).
- HABCs are derived by using approximations of exact non-local boundary operators or exact exterior solutions, which result in local equations to be solved on the artificial boundary. The accuracy of each HABC depends on the kind and the order of approximation. Early HABCs have been written with high-order partial derivatives in time and space [Engquist and Majda, 1977, 1979; Higdon, 1986], where the order of derivatives is as high as the order of approximation. Such formulations have been successfully implemented in time-harmonic contexts [Givoli et al., 1997; Schmidt et al., 2015], but their applicability is limited to low orders because of the high-order derivatives that must be discretized. As an alternative strategy, Collino [Collino, 1993a,b] proposed a formulation with only low-order derivatives and auxiliary fields defined on the boundary, enabling the use of HABCs with high orders. This strategy has be applied to derive transient HABCs in different physical and numerical contexts [Baffet et al., 2012; Bécache et al., 2010; Diaz and Joly, 2005; Givoli, 2001; Givoli and Neta, 2003; Givoli et al., 2006; Guddati and Lim, 2006; Guddati and Tassoulas, 2000; Hagstrom and Warburton, 2004, 2009; Hagstrom et al., 2010; Rabinovich et al., 2011].

In the last two decades, PMLs have clearly received much more attention than HABCs. One reason is that the PMLs are easier to implement than HABCs. While HABCs require specific resolution procedures, with often cumbersome treatments of the corners of the computational domain, the PMLs can be rather easily implemented in existing computational codes with straightforward treatments at the corners. To the best of our knowledge, HABCs have never been implemented with discontinuous Galerkin schemes in three dimensions, by contrast with PMLs (see *e.g.* [Lähivaara and Huttunen, 2010; Lu et al., 2004; Modave et al., 2016a]). Nevertheless, the accuracy of PMLs strongly depends on the discretization and the choice of parameters that must be carefully chosen (see *e.g.* [Bermúdez et al., 2007; Collino and Monk, 1998; Modave et al., 2014]), while *a priori* analyses can be used to optimize HABCs [Hagstrom and Warburton, 2009]. Recently, the double absorbing boundary (DAB) [Baffet et al., 2014; Hagstrom et al., 2014; LaGrone and Hagstrom, 2016] has been introduced to simplify the treatment of corners in HABC procedures. The DAB technique relies on an extension of the domain with a thin layer where the auxiliary fields of the HABCs are defined. Additional discussions and comparisons of boundary treatments can be found in [Givoli, 2004; Hagstrom, 1999; Rabinovich et al., 2010]

In this paper, we propose a computational approach that combines a family of local HABCs with a nodal discontinuous Galerkin finite element method for 3D acoustic wave simulations in cuboidal domains. The HABCs are built in a classical way by using Padé approximations of a square root in the exact boundary condition. By choosing a specific representation for these HABCs, we derive edge and corner compatibility conditions which are easier to use than those obtained with existing Padé-type HABCs. The proposed formulations overcome some inconsistencies that appear when deriving discontinuous Galerkin schemes. Our

strategy is inspired by the works of Hagstrom and Warburton [Hagstrom and Warburton, 2004, 2009], where 2D discontinuous Galerkin results are presented with HABCs.

In order to demonstrate the applicability of HABC techniques with discontinuous Galerkin methods, we describe a GPU-accelerated implementation and present the results of a 3D realistic benchmark. The complete boundary procedure requires a multidimensional solver with equations posed in the volume, on the faces, the edges and the corners of the domain. We use elaborate implementation techniques in order to improve the computational performance, while keeping the compatibility of the final implementation with implementations in the literature [Chan et al., 2016b; Modave et al., 2015]. In order to deal with large scale problems, the boundary procedure could be coupled naturally with multi-rate time stepping schemes, discontinuous Galerkin schemes for hybrid meshes and computational strategies for computations on GPU clusters.

This paper is organized as follows. In section 2, the HABCs are presented and edge/corner compatibility conditions are derived for both the wave equation and the pressure-velocity system. Section 3 is dedicated to numerical schemes and implementation strategies. We describe the discontinuous Galerkin finite element scheme, the low-storage Runge-Kutta scheme, and key aspects of the GPU-accelerated implementation. In section 4, we present three-dimensional numerical results which validate the formulation and demonstrate the computational performance and the applicability of the approach.

2 High-order absorbing boundary treatment for cuboidal domain

In this section, we derive high-order absorbing boundary conditions (HABCs) with compatibility conditions for edges and corners to simulate the propagation of waves in the infinite space with a cuboidal computational domain. We aim at coupling these conditions with a numerical scheme based on the pressure-velocity system

$$\frac{\partial p}{\partial t} + \rho c^2 \nabla \cdot \mathbf{u} = 0, \tag{1a}$$

$$\rho \frac{\partial \mathbf{u}}{\partial t} + \nabla p = 0, \tag{1b}$$

where $p(t, \mathbf{x})$ is the pressure field, $\mathbf{u}(t, \mathbf{x})$ is the velocity field, ρ is the density and c is the phase velocity. In this section, ρ and c are assumed to be constant. When deriving the conditions, it is however more convenient to work with the wave equation

$$\frac{\partial^2 p}{\partial t^2} - c^2 \Delta p = 0. \tag{2}$$

The pressure-velocity system can be recovered by introducing the velocity field governed by equation (1b) and integrating the wave equation.

We first derive HABCs for a semi-infinite domain with a planar boundary (section 2.1). When using the obtained HABCs on all the faces of a cuboidal domain, a special treatment must be applied to the edges and the corners. A treatment based on accuracy-preserving compatibility conditions is proposed in section 2.2. The obtained equations of both the HABCs and the compatibility conditions involve second-order partial derivatives. In section 2.3, we derive equivalent formulations written using only first-order partial derivatives, which can quite naturally be coupled with explicit time-stepping schemes and spatial schemes based on the pressure-velocity system.

2.1 High-order absorbing boundary conditions for planar boundary

Let us consider the half-space problem defined on the domain $\Omega = \{\mathbf{x} \in \mathbb{R}^3 : x < 0\}$ with the planar boundary $\Gamma = \{\mathbf{x} \in \mathbb{R}^3 : x = 0\}$, where x is the coordinate in the Cartesian direction $\mathbf{e}_x$. We seek a non-reflective

boundary condition to prescribe on Γ. For convenience, the transverse component of the position is denoted $\mathbf{y}$, such that $\mathbf{x} = (x, \mathbf{y})$.

The exact non-reflective boundary condition of the half-space problem is well-known. In the Fourier domain, this condition reads (see *e.g.* [Engquist and Majda, 1977; Hagstrom, 1999])

$$c\frac{d\tilde{p}}{dx} + \imath\omega\sqrt{1-(c\|\boldsymbol{\xi}\|/\omega)^2}\,\tilde{p} = 0,$$

where ω and $\boldsymbol{\xi}$ are the dual Fourier variables to t and $\mathbf{y}$, respectively, and $\tilde{p}(\omega, x, \boldsymbol{\xi})$ is the Fourier transform of $p(t, x, \mathbf{y})$. We take the convention $\partial_t \leftrightarrow \imath\omega$ and $\nabla_\perp \leftrightarrow \imath\boldsymbol{\xi}$ for the Fourier transform, where $\nabla_\perp$ is the transverse component of the operator ∇ in the plan tangent to the direction $\mathbf{e}_x$. The time-space version of the condition is obtained using an inverse Fourier transform, which gives

$$c\frac{dp}{dx} + \iint \imath\omega\sqrt{1-(c\|\boldsymbol{\xi}\|/\omega)^2}\,\tilde{p}\,e^{\imath(\omega t+\boldsymbol{\xi}\cdot\mathbf{y})}\,d\omega\,d\boldsymbol{\xi} = 0. \tag{3}$$

Unfortunately, this condition is non-local in both time and space because of the square root, which makes it an impractical boundary treatment. Using notations borrowed from the pseudo-differential theory, the condition can be rewritten as

$$\mathcal{B}^x p = 0, \tag{4}$$

with the pseudo-differential operator

$$\mathcal{B}^x \equiv c\,\partial_x + \partial_t\sqrt{1-c^2\Delta_\perp^x/\partial_{tt}}, \tag{5}$$

where $\Delta_\perp^x \equiv \Delta - \partial_{xx}$ is the Laplace-Beltrami operator defined in the plan tangent to the direction $\mathbf{e}_x$.

Local absorbing boundary conditions are classically obtained by approximating the square root $\sqrt{1+X}$ to localize the operator (5). The features of the obtained conditions depend on the approximation that is used for the square root. In their seminal work, Engquist and Majda [Engquist and Majda, 1977] showed that Padé approximations lead to stable conditions, while polynomial approximations based on Taylor expansions can lead to unstable conditions. Unfortunately, while the Padé-type conditions simulate outgoing traveling waves with an arbitrarily high accuracy, they are inefficient for grazing and evanescent waves. Higdon [Higdon, 1986] proposed a rational approximation of the square root that increases accuracy for grazing waves. Other approximations that address one or both drawbacks of Padé approximations have been used to derive absorbing boundary conditions and one-way wave equations (see *e.g.* [Antoine et al., 2006; Asvadurov et al., 2003; Hagstrom and Warburton, 2009; Hagstrom et al., 2010; Halpern and Trefethen, 1988; Ingerman et al., 2000; Kechroud et al., 2005; Lu, 1998; Milinazzo et al., 1997]). In this work, we restrict ourselves to the Padé approximation, which represents a fundamental case. The boundary treatment and the computational procedure described hereafter will be extended to more accurate approximations in the future.

We consider the $(2N+1)^{\text{th}}$-order Padé approximation of the square root $\sqrt{1+X}$ written as the rational function [Bamberger et al., 1988; Guan-Quan, 1985; Lu, 1998]

$$f_N(X) = 1 + \frac{2}{M}\sum_{n=1}^{N}\frac{a_n X}{1+b_n X}, \tag{6}$$

$$= 1 + \frac{2}{M}\sum_{n=1}^{N} c_n\left(1 - \frac{1+c_n}{1+c_n+X}\right), \tag{7}$$

where $a_n = \sin^2(n\pi/M)$, $b_n = \cos^2(n\pi/M)$, $c_n = \tan^2(n\pi/M)$ and $M = 2N+1$. Using the formula (7) to approximate the square root in the exact boundary operator (5) gives to the approximate boundary condition

$$\mathcal{L}^x p = 0, \tag{8}$$

with the pseudo-differential operator

$$\mathcal{L}^x \equiv \partial_t + c\,\partial_x + \frac{2}{M}\sum_{n=1}^{N} c_n\,\partial_t\left(1 - \frac{(1+c_n)\partial_{tt}}{(1+c_n)\partial_{tt} - c^2\Delta_\perp^x}\right).$$

Following a strategy proposed by Collino [Collino, 1993a,b], we introduce N auxiliary fields p_n defined on the boundary Γ as

$$p_n \equiv \mathcal{M}_n^x\, p, \qquad \text{for } n = 1, \ldots, N,$$

with

$$\mathcal{M}_n^x \equiv \frac{(1+c_n)\partial_{tt}}{(1+c_n)\partial_{tt} - c^2\Delta_\perp^x}, \qquad \text{for } n = 1, \ldots, N. \tag{9}$$

We can then explicitly rewrite the boundary condition (8) as

$$\partial_t p + c\,\partial_x p = \frac{2}{M}\sum_{n=1}^{N} c_n \partial_t (p_n - p), \tag{10}$$

where the auxiliary fields are governed by

$$(1+c_n)\,\partial_{tt}\,(p_n - p) - c^2\Delta_\perp^x p_n = 0, \qquad \text{for } n = 1, \ldots, N. \tag{11}$$

The boundary condition is local and requires the computation of the N auxiliary equations on the boundary Γ. Increasing the order of the Padé approximation increases the number of auxiliary equations and the computational cost, but it also improves the accuracy for outgoing traveling waves [Engquist and Majda, 1977]. Since equation (10) can theoretically be used at any order of accuracy, this condition belongs to the family of HABCs.

We note that rational approximations of the square root can be written in several ways, leading to different representations of the boundary conditions. Engquist and Majda [Engquist and Majda, 1977] directly wrote the Padé-type boundary conditions using high-order partial derivatives in time and space, without any additional fields. However, such formulations require accurate numerical schemes for high-order derivatives, which limits the applicability with standard finite difference and finite element discretizations. Collino [Collino, 1993a,b] used the rational representation (6) and obtained HABCs with auxiliary equations very close to equations (10)-(11). Alternatively, rational functions written as continued fractions lead to HABCs with auxiliary fields governed by coupled equations (see *e.g.* [Asvadurov et al., 2003; Givoli and Neta, 2003; Guddati and Tassoulas, 2000; Hagstrom and Warburton, 2004]). In this work, we choose the representation (7) because it leads to accurate compatibility conditions that are cheaper to use and easier to incorporate in a discontinuous Galerkin framework than those obtained in previous works. Both aspects are discussed later in the text.

2.2 Compatibility conditions at edges and corners

We now extend the boundary treatment to the borders of a cuboidal domain $\Omega = [-L^x, L^x] \times [-L^y, L^y] \times [-L^z, L^z]$ to simulate the infinite space $\mathbb{R}^3$. The initial conditions and any source are assumed to be compactly supported inside the domain Ω. Under this assumption, the exact boundary condition (4) and its approximation (8) can be prescribed on the planes containing each of the six faces of the domain (figure 1a). In practice, we would like to prescribe the boundary condition (10) only on the faces (figure 1b). However, computing the auxiliary equations (11) only on the faces requires boundary conditions for the auxiliary fields on the borders of each face, which are on the edges of the domain.

To derive such conditions, we have adapted a strategy proposed by Hagstrom and Warburton [Hagstrom and Warburton, 2004]. This strategy is based on purely algebraic manipulations of the equations which give

(a) HABC defined on planes

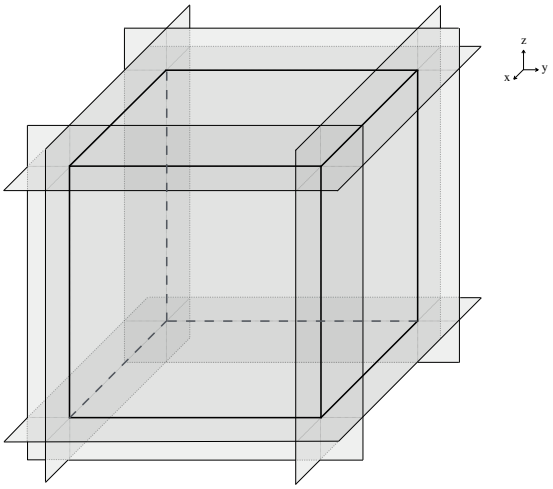

(b) HABC defined on faces

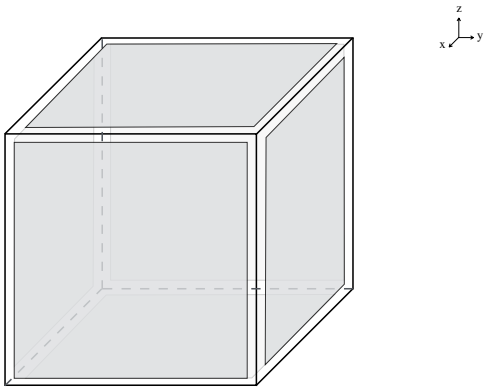

Figure 1: Illustration of HABCs defined on the planes containing the faces of the cuboidal domain (a) or defined only on the faces (b). When defining HABCs only on the faces, supplementary boundary conditions must be prescribed for the auxiliary variables on the borders of each face.

compatibility conditions on the edges. These conditions preserve the accuracy of the solution, and require the computation of supplementary auxiliary fields governed by auxiliary equations on the edges. Proceeding in a hierarchical fashion, these new equations require boundary conditions on the borders of each edge, which are at the corners of the domain. Algebraic manipulations of these equations similarly provide compatibility conditions with supplementary auxiliary fields at the corners.

In this section, we derive compatibility conditions for the edges belonging to the lines $(x, y) = (L^x, L^y)$, $(x, z) = (L^x, L^z)$ and $(y, z) = (L^y, L^z)$ and for the corner $(x, y, z) = (L^x, L^y, L^z)$. The conditions for the other edges and corners can be obtained straightforwardly by symmetry. For the sake of clarity, we first define the fields as

$$p_{ijk} \equiv \mathcal{M}_i^x \mathcal{M}_j^y \mathcal{M}_k^z \, p, \qquad \text{for } i, j, k = 0, \dots, N, \tag{12}$$

where the operators $\mathcal{M}_i^x$, $\mathcal{M}_j^y$, $\mathcal{M}_k^z$ are defined using equation (9) for $i, j, k > 0$ and $\mathcal{M}_0^x = \mathcal{M}_0^y = \mathcal{M}_0^z = 1$. The fields with zero, one, two and three non-zero subscripts are computed on the volume, the faces, the edges and the corners, respectively. The field p_{000} corresponds to the pressure field p. The fields p_{i00}, p_{0j0}, p_{00k}, with $i, j, k = 1, \dots, N$, are defined on the faces $x = L^x$, $y = L^y$, $z = L^z$, respectively. The fields p_{ij0}, p_{i0k}, p_{0jk}, with $i, j, k = 1, \dots, N$, are defined on the edges $(x, y) = (L^x, L^y)$, $(x, z) = (L^x, L^z)$, $(y, z) = (L^y, L^z)$, respectively. The fields p_{ijk}, with $i, j, k = 1, \dots, N$, are defined on the corner $(x, y, z) = (L^x, L^y, L^z)$. Therefore, there are N, N^2 and N^3 auxiliary fields per face, edge and corner, respectively (figure 2). Note that, since the initial conditions on p are assumed to be compactly supported inside the domain, the initial conditions on the auxiliary fields are equal to 0.

The auxiliary fields have two properties that are used when deriving the compatibility conditions. First, all the auxiliary fields satisfy the wave equation,

$$\partial_{tt} p_{ijk} - c^2 \Delta p_{ijk} = 0, \qquad \text{for } i, j, k = 0, \dots, N. \tag{13}$$

Indeed, since the pressure field is in the null space of the wave operator $\left[\partial_{tt} - c^2\Delta\right]$, the auxiliary fields defined by equation (12) are also in this null space. Then, the auxiliary fields are related by the following

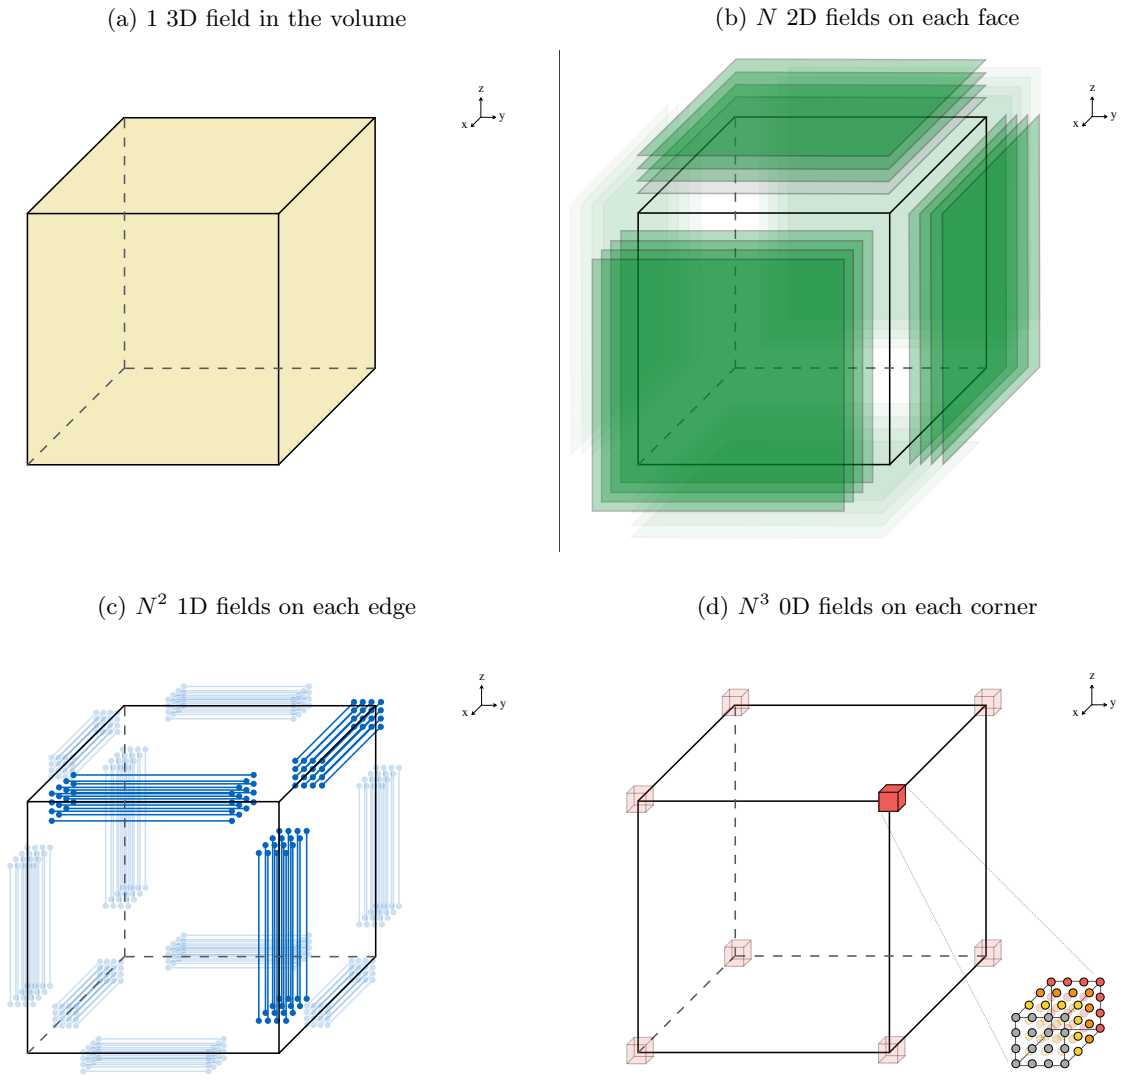

Figure 2: Illustration of the fields defined in the volume (a), on the faces (b), on the edges (c) and on the corners (d) of a cuboidal domain. These pictures show the case $N = 4$. The equations corresponding to the faces/edges/corner in darker color are explicitly derived in the text. The equations corresponding to the other faces/edges/corners can straightforwardly be obtained by symmetry.

relations

$$p_{ijk} = \mathcal{M}_i^x \, p_{0jk}, \qquad \text{for } i,j,k = 0,\dots,N, \tag{14a}$$
$$p_{ijk} = \mathcal{M}_j^y \, p_{i0k}, \qquad \text{for } i,j,k = 0,\dots,N, \tag{14b}$$
$$p_{ijk} = \mathcal{M}_k^z \, p_{ij0}, \qquad \text{for } i,j,k = 0,\dots,N. \tag{14c}$$

These relations are trivially obtained by using definition (12) and observing that the operators $\mathcal{M}_i^x$, $\mathcal{M}_j^y$, $\mathcal{M}_k^z$ commute.

Boundary conditions on faces and 2D relations

We write the boundary conditions for the 3D field on the faces $x = L^x$, $y = L^y$ and $z = L^z$. On these faces, the field p_{000} satisfies HABCs corresponding to half spaces,

$$\begin{aligned} \mathcal{L}^x p_{000} &= 0, \qquad \text{for } x = L^x, \\ \mathcal{L}^y p_{000} &= 0, \qquad \text{for } y = L^y, \\ \mathcal{L}^z p_{000} &= 0, \qquad \text{for } z = L^z, \end{aligned}$$

which can be rewritten

$$\partial_t p_{000} + c\,\partial_x p_{000} = \frac{2}{M}\sum_{i=1}^{N} c_i \partial_t (p_{i00} - p_{000}), \qquad \text{for } x = L^x, \tag{15a}$$

$$\partial_t p_{000} + c\,\partial_y p_{000} = \frac{2}{M}\sum_{j=1}^{N} c_j \partial_t (p_{0j0} - p_{000}), \qquad \text{for } y = L^y, \tag{15b}$$

$$\partial_t p_{000} + c\,\partial_z p_{000} = \frac{2}{M}\sum_{k=1}^{N} c_k \partial_t (p_{00k} - p_{000}), \qquad \text{for } z = L^z. \tag{15c}$$

Using the definition (12), we have the 2D relations

$$(1 + c_i)\partial_{tt}(p_{i00} - p_{000}) - c^2(\partial_{yy} + \partial_{zz})p_{i00} = 0, \qquad \text{for } x = L^x, \tag{16a}$$

$$(1 + c_j)\partial_{tt}(p_{0j0} - p_{000}) - c^2(\partial_{xx} + \partial_{zz})p_{0j0} = 0, \qquad \text{for } y = L^y, \tag{16b}$$

$$(1 + c_k)\partial_{tt}(p_{00k} - p_{000}) - c^2(\partial_{xx} + \partial_{yy})p_{00k} = 0, \qquad \text{for } z = L^z, \tag{16c}$$

for $i, j, k = 1, \ldots, N$.

Boundary conditions on edges and 1D relations

On the borders of each face (*i.e.* on the edges of the domain), the 2D fields need boundary conditions. We derive here the compatibility conditions for the 2D fields on the edges $(x, y) = (L^x, L^y)$, $(x, z) = (L^x, L^z)$ and $(y, z) = (L^y, L^z)$.

Because p_{000} belongs to the null space of $\mathcal{L}^x$ and $\mathcal{L}^y$ on the edge $(x, y) = (L^x, L^y)$ and because these operators commute with $\mathcal{M}^x$ and $\mathcal{M}^y$, the 2D fields p_{0j0} and p_{i00} also belong to the null space of $\mathcal{L}^x$ and $\mathcal{L}^y$ seen their definition. This result gives boundary conditions for the 2D fields on the edge $(x, y) = (L^x, L^y)$. Using similar reasoning, we obtain boundary conditions on the other edges. We then have

$$\begin{aligned} \mathcal{L}^x p_{0j0} &= 0, \qquad \text{for } (x, y) = (L^x, L^y), \\ \mathcal{L}^y p_{i00} &= 0, \qquad \text{for } (x, y) = (L^x, L^y), \\ \mathcal{L}^x p_{00k} &= 0, \qquad \text{for } (x, z) = (L^x, L^z), \\ \mathcal{L}^z p_{i00} &= 0, \qquad \text{for } (x, z) = (L^x, L^z), \\ \mathcal{L}^y p_{00k} &= 0, \qquad \text{for } (y, x) = (L^y, L^z), \\ \mathcal{L}^z p_{0j0} &= 0, \qquad \text{for } (y, x) = (L^y, L^z), \end{aligned}$$

which can be rewritten

$$\partial_t p_{0j0} + c\,\partial_x p_{0j0} = \frac{2}{M}\sum_{i=1}^{N} c_i \partial_t (p_{ij0} - p_{0j0}), \qquad \text{for } (x,y) = (L^x, L^y), \tag{17a}$$

$$\partial_t p_{i00} + c\,\partial_y p_{i00} = \frac{2}{M}\sum_{j=1}^{N} c_j \partial_t (p_{ij0} - p_{i00}), \qquad \text{for } (x,y) = (L^x, L^y), \tag{17b}$$

$$\partial_t p_{00k} + c\,\partial_x p_{00k} = \frac{2}{M}\sum_{i=1}^{N} c_i \partial_t (p_{i0k} - p_{00k}), \qquad \text{for } (x,z) = (L^x, L^z), \tag{17c}$$

$$\partial_t p_{i00} + c\,\partial_z p_{i00} = \frac{2}{M}\sum_{k=1}^{N} c_k \partial_t (p_{i0k} - p_{i00}), \qquad \text{for } (x,z) = (L^x, L^z), \tag{17d}$$

$$\partial_t p_{00k} + c\,\partial_y p_{00k} = \frac{2}{M}\sum_{j=1}^{N} c_j \partial_t (p_{0jk} - p_{00k}), \qquad \text{for } (y,z) = (L^y, L^z), \tag{17e}$$

$$\partial_t p_{0j0} + c\,\partial_z p_{0j0} = \frac{2}{M}\sum_{k=1}^{N} c_k \partial_t (p_{0jk} - p_{0j0}), \qquad \text{for } (y,z) = (L^y, L^z), \tag{17f}$$

for $i, j, k = 1, \ldots, N$.

We next derive governing equations for the 1D fields on the edges. Using equations (14a)-(14c) to connect the 1D and the 2D fields, one has

$$\partial_{tt}(1 + c_i)(p_{ij0} - p_{0j0}) - c^2(\partial_{yy} + \partial_{zz})p_{ij0} = 0, \qquad \text{for } (x,y) = (L^x, L^y), \tag{18a}$$

$$\partial_{tt}(1 + c_j)(p_{ij0} - p_{i00}) - c^2(\partial_{xx} + \partial_{zz})p_{ij0} = 0, \qquad \text{for } (x,y) = (L^x, L^y), \tag{18b}$$

$$\partial_{tt}(1 + c_i)(p_{i0k} - p_{00k}) - c^2(\partial_{yy} + \partial_{zz})p_{i0k} = 0, \qquad \text{for } (x,z) = (L^x, L^z), \tag{18c}$$

$$\partial_{tt}(1 + c_k)(p_{i0k} - p_{i00}) - c^2(\partial_{xx} + \partial_{yy})p_{i0k} = 0, \qquad \text{for } (x,z) = (L^x, L^z), \tag{18d}$$

$$\partial_{tt}(1 + c_j)(p_{0jk} - p_{00k}) - c^2(\partial_{xx} + \partial_{zz})p_{0jk} = 0, \qquad \text{for } (y,z) = (L^y, L^z), \tag{18e}$$

$$\partial_{tt}(1 + c_k)(p_{0jk} - p_{0j0}) - c^2(\partial_{xx} + \partial_{yy})p_{0jk} = 0, \qquad \text{for } (y,z) = (L^y, L^z), \tag{18f}$$

for $i, j, k = 1, \ldots, N$. Unfortunately, these relations involve spatial derivatives that cannot be computed for 1D fields defined only on the edges. For instance, only derivatives with respect to z can be computed on the edge $(x, y) = (L^x, L^y)$, while equations (18a) and (18b) involve derivatives with respect to y and x, respectively. We then manipulate equations (18a)-(18f) to eliminate such inadmissible derivatives. Adding the equations corresponding to each edge and using equation (13) give

$$\partial_{tt}\left[(1 + c_i + c_j)p_{ij0} - (1 + c_i)p_{0j0} - (1 + c_j)p_{i00}\right] - c^2\partial_{zz}p_{ij0} = 0, \qquad \text{for } (x,y) = (L^x, L^y), \tag{19a}$$

$$\partial_{tt}\left[(1 + c_i + c_k)p_{i0k} - (1 + c_i)p_{00k} - (1 + c_k)p_{i00}\right] - c^2\partial_{yy}p_{i0k} = 0, \qquad \text{for } (x,z) = (L^x, L^z), \tag{19b}$$

$$\partial_{tt}\left[(1 + c_j + c_k)p_{0jk} - (1 + c_j)p_{00k} - (1 + c_k)p_{0j0}\right] - c^2\partial_{xx}p_{0jk} = 0, \qquad \text{for } (y,z) = (L^y, L^z), \tag{19c}$$

for $i, j, k = 1, \ldots, N$. These 1D relations involve only spatial derivatives which are well-defined on edges.

Boundary conditions at corners and 0D relations

On the borders of each edge (*i.e.* at the corners of the domain), the 1D field equations require boundary conditions. We derive the compatibility conditions for the 1D fields at the corner $(x, y, z) = (L^x, L^y, L^z)$ by using a similar strategy than for the edges.

Because p_{000} belongs to the null space of $\mathcal{L}^x$, $\mathcal{L}^y$ and $\mathcal{L}^z$ and because these operators commute with $\mathcal{M}^x$, $\mathcal{M}^y$ and $\mathcal{M}^z$, the 1D fields p_{0jk}, p_{i0k} and p_{ij0} also belong to the null space of $\mathcal{L}^x$, $\mathcal{L}^y$ and $\mathcal{L}^z$ by definition

(12). At the corner $(x, y, z) = (L^x, L^y, L^z)$, we then have

$$\begin{aligned}\mathcal{L}^x p_{0jk} &= 0,\\ \mathcal{L}^y p_{i0k} &= 0,\\ \mathcal{L}^z p_{ij0} &= 0,\end{aligned}$$

which can be rewritten

$$\partial_t p_{0jk} + c\,\partial_x p_{0jk} = \frac{2}{M}\sum_{i=1}^{N} c_i \partial_t (p_{ijk} - p_{0jk}), \tag{20a}$$

$$\partial_t p_{i0k} + c\,\partial_y p_{i0k} = \frac{2}{M}\sum_{j=1}^{N} c_j \partial_t (p_{ijk} - p_{i0k}), \tag{20b}$$

$$\partial_t p_{ij0} + c\,\partial_z p_{ij0} = \frac{2}{M}\sum_{k=1}^{N} c_k \partial_t (p_{ijk} - p_{ij0}), \tag{20c}$$

for $i, j, k = 1, \ldots, N$.

We now derive relations for the 0D auxiliary fields at the corner. Using equations (14a)-(14c) to connect the 0D and the 1D fields, one has

$$(1 + c_i)\partial_{tt}(p_{ijk} - p_{0jk}) - c^2(\partial_{yy} + \partial_{zz})p_{ijk} = 0, \tag{21a}$$

$$(1 + c_j)\partial_{tt}(p_{ijk} - p_{i0k}) - c^2(\partial_{xx} + \partial_{zz})p_{ijk} = 0, \tag{21b}$$

$$(1 + c_k)\partial_{tt}(p_{ijk} - p_{ij0}) - c^2(\partial_{xx} + \partial_{yy})p_{ijk} = 0, \tag{21c}$$

for $i, j, k = 1, \ldots, N$. Unfortunately, these relations again involve spatial derivatives which are inadmissible if the 0D fields are defined only on the corner. Again, we manipulate the equations to eliminate these derivatives. Adding equations (21a)–(21c), removing the spatial derivatives by using equation (13) and integrating in time give

$$(1 + c_i + c_j + c_k)p_{ijk} - (1 + c_i)p_{0jk} - (1 + c_j)p_{i0k} - (1 + c_k)p_{ij0} = 0, \tag{22}$$

for $i, j, k = 1, \ldots, N$. These 0D relations do not involve any derivatives, and they can be used to remove the 0D fields p_{ijk} from the boundary conditions (20a)-(20c).

Summary and extension to mixed boundary conditions

The boundary treatment presented in this section consists of a multidimensional solver with equations to solve in the volume, on the faces, on the edges and at the corners of the domain. The formulation is summarized in algorithm 1. By the assumption of the compact support of the initial condition inside the domain, the initial conditions for the auxiliary variable fields are all zero.

We briefly address the case where a homogeneous boundary condition is prescribed on one or more faces of the computational domain. In exploration geophysics, for instance, the computational cuboidal domain must represent the underground structure. A HABC can be used on the lateral and bottom faces of the domain, while the so-called *free-surface* boundary condition (which is $p = 0$ in the acoustic model) must be prescribed on the upper face to simulate the Earth's surface.

A homogeneous boundary condition on the 3D fields is straightforwardly incorporated into the boundary procedure by using the same condition on the auxiliary 2D and 1D fields. Indeed, let us consider the homogeneous Dirichlet condition $p = 0$ on one face of the domain and a HABC on the adjacent faces. Seen definition (12), the 2D auxiliary fields of the adjacent faces must be set to zero on the edges if $p = 0$. Therefore, no auxiliary field must be computed on edges and corners belonging to faces where a homogeneous boundary condition is prescribed.

Algorithm 1: Multidimensional solver for the wave equation with HABC and edges/corners compatibility. The equation numbers correspond to face $x = L^x$, edge $(x, y) = (L^x, L^y)$ and corner $(x, y, z) = (L^x, L^y, L^z)$.

3D solver (on the volume):

- compute the 3D pressure field governed by the classical wave equation with initial conditions and/or sources which are compactly supported inside the domain;

2D solver (on the faces):

- prescribe one boundary condition using the first-order differential equation (15a);
- compute N 2D auxiliary fields governed by the second-order differential equations (16a);

1D solver (on the edges):

- prescribe $2N$ boundary conditions (N for each adjacent face) using the first-order differential equations (17a)-(17b);
- compute N^2 1D auxiliary fields governed by the second-order differential equations (19a);

0D solver (at the corners):

- prescribe $3N^2$ boundary conditions (N^2 for each adjacent edge) using the first-order differential equations (20a)-(20c), where N^3 0D auxiliary fields are defined by the algebraic relations (22).

2.3 Formulation with first-order systems and characteristic variables

Since all the auxiliary fields $p_{ijk}(t, \mathbf{x})$ satisfy the wave equation (13), we can define auxiliary velocities $\mathbf{u}_{ijk}(t, \mathbf{x})$ such that each pair $(p_{ijk}, \mathbf{u}_{ijk})$ satisfies the pressure-velocity system

$$\partial_t p_{ijk} + \rho c^2 \, \nabla \cdot \mathbf{u}_{ijk} = 0, \tag{23a}$$

$$\rho \, \partial_t \mathbf{u}_{ijk} + \nabla p_{ijk} = 0, \tag{23b}$$

with $i, j, k = 0, \ldots, N$. Introducing these auxiliary velocities and using integration in time on the equations derived in the previous section give a HABC formulation with only algebraic relations as boundary conditions and first-order differential equations as governing equations for the auxiliary fields:

- on the face $x = L^x$, the boundary condition becomes

$$p_{000} - \rho c u_{000} = \frac{2}{M} \sum_{i=1}^{N} c_i \, (p_{i00} - p_{000}), \tag{24}$$

and the 2D fields are governed by

$$\begin{aligned} (1 + c_i)\partial_t p_{i00} + \rho c^2 \big(\partial_y v_{i00} + \partial_z w_{i00}\big) &= (1 + c_i)\partial_t p_{000}, \\ \rho \partial_t v_{i00} + \partial_y p_{i00} &= 0, \\ \rho \partial_t w_{i00} + \partial_z p_{i00} &= 0, \end{aligned} \tag{25}$$

for $i = 1, \ldots, N$;

- on the edge $(x, y) = (L^x, L^y)$, the boundary conditions become

$$p_{0j0} - \rho c u_{0j0} = \frac{2}{M} \sum_{i=1}^{N} c_i (p_{ij0} - p_{0j0}),$$
$$p_{i00} - \rho c v_{i00} = \frac{2}{M} \sum_{j=1}^{N} c_j (p_{ij0} - p_{i00}),$$

and the 1D fields are governed by

$$\begin{aligned} (1 + c_i + c_j)\partial_t p_{ij0} + \rho c^2 \partial_z w_{ij0} &= (1 + c_j)\partial_t p_{i00} + (1 + c_i)\partial_t p_{0j0} \\ \rho \partial_t w_{ij0} + \partial_z p_{ij0} &= 0, \end{aligned} \tag{26}$$

for $i, j = 1, \ldots, N$;

- at the corner $(x, y, z) = (L^x, L^y, L^z)$, the boundary conditions become

$$p_{0jk} - \rho c u_{0jk} = \frac{2}{M} \sum_{i=1}^{N} c_i \big((1 - C_{ijk}) p_{0jk} + C_{jki} p_{i0k} + C_{kij} p_{ij0} \big),$$
$$p_{i0k} - \rho c v_{i0k} = \frac{2}{M} \sum_{j=1}^{N} c_j \big(C_{ijk} p_{0jk} + (1 - C_{jki}) p_{i0k} + C_{kij} p_{ij0} \big),$$
$$p_{ij0} - \rho c v_{ij0} = \frac{2}{M} \sum_{k=1}^{N} c_k \big(C_{ijk} p_{0jk} + C_{jki} p_{i0k} + (1 - C_{kij}) p_{ij0} \big),$$

with

$$C_{ijk} = \frac{1 + c_i}{1 + c_i + c_j + c_k},$$

for $i, j, k = 1, \ldots, N$.

Similar relations can be obtained for the other faces, edges and corners.

Implementing this formulation with an explicit time-stepping scheme is not straightforward because the differential equations governing the pressure fields are coupled by terms with time derivatives. Indeed, the time derivatives of both 3D pressure and 2D pressures appear in equation (25), and the time derivatives of both 2D pressures and 1D pressures appear in equation (26). In addition, all the fields are also coupled through the algebraic relations in the different levels.

In order to efficiently implement the boundary procedure with an explicit time-stepping scheme and a discontinuous Galerkin method, we reformulate the boundary conditions using characteristic variables (as in [Hagstrom and Warburton, 2004]) and introducing temporary variables. In the directions $\mathbf{e}_x$, $\mathbf{e}_y$, $\mathbf{e}_z$, the characteristic variables read

$$r^{-x}_{ijk} = p_{ijk} - \rho c\, u_{ijk}, \qquad r^{+x}_{ijk} = p_{ijk} + \rho c\, u_{ijk}, \tag{27a}$$
$$r^{-y}_{ijk} = p_{ijk} - \rho c\, v_{ijk}, \qquad r^{+y}_{ijk} = p_{ijk} + \rho c\, v_{ijk}, \tag{27b}$$
$$r^{-z}_{ijk} = p_{ijk} - \rho c\, w_{ijk}, \qquad r^{+z}_{ijk} = p_{ijk} + \rho c\, w_{ijk}, \tag{27c}$$

where r^{-x}_{ijk}, r^{-y}_{ijk}, r^{-z}_{ijk} and r^{+x}_{ijk}, r^{+y}_{ijk}, r^{+z}_{ijk} contain information traveling downwardly and upwardly, respectively, along the direction $\mathbf{e}_x$, $\mathbf{e}_y$, $\mathbf{e}_z$. The comprehensive boundary formulation reads

- on the face $x = L^x$, temporary fields $p^\star_{i00}$ are defined as

$$p^\star_{i00} = p_{i00} - (r^{-x}_{000} + r^{+x}_{000}), \tag{28}$$

the boundary condition gives

$$r^{-x}_{000} = \frac{1}{M}\sum_{i=1}^{N} c_i\, p^\star_{i00}, \tag{29}$$

and the 2D relations becomes

$$(1+c_i)\ \partial_t p^\star_{i00} + \rho c^2\big(\partial_y v_{i00} + \partial_z w_{i00}\big) = 0, \tag{30a}$$
$$\rho\partial_t v_{i00} + \partial_y p_{i00} = 0, \tag{30b}$$
$$\rho\partial_t w_{i00} + \partial_z p_{i00} = 0, \tag{30c}$$

for $i = 1, \ldots, N$;

- on the edge $(x, y) = (L^x, L^y)$, temporary fields $p^\star_{ij0}$, with $i, j = 1, \ldots, N$, are defined as

$$p^\star_{ij0} = p_{ij0} - \big(C_{ji}(r^{-y}_{i00} + r^{+y}_{i00}) + C_{ij}(r^{-x}_{0j0} + r^{+x}_{0j0})\big), \tag{31}$$

with

$$C_{ij} = \frac{1+c_i}{1+c_i+c_j},$$

the boundary conditions give

$$r^{-y}_{i00} = \frac{1}{M}\sum_{j=1}^{N} c_j\ \big(p^\star_{ij0} + (C_{ji}-1)\,(r^{-y}_{i00} + r^{+y}_{i00}) + C_{ij}\,(r^{-x}_{0j0} + r^{+x}_{0j0})\big), \tag{32a}$$

$$r^{-x}_{0j0} = \frac{1}{M}\sum_{i=1}^{N} c_i\ \big(p^\star_{ij0} + C_{ji}\,(r^{-y}_{i00} + r^{+y}_{i00}) + (C_{ij}-1)\,(r^{-x}_{0j0} + r^{+x}_{0j0})\big), \tag{32b}$$

and the 1D relations becomes

$$(1+c_i+c_j)\ \partial_t p^\star_{ij0} + \rho c^2 \partial_z w_{ij0} = 0, \tag{33a}$$
$$\rho\partial_t w_{ij0} + \partial_z p_{ij0} = 0, \tag{33b}$$

for $i, j = 1, \ldots, N$;

- at the corner $(x, y, z) = (L^x, L^y, L^z)$, the boundary conditions give

$$r^{-x}_{0jk} = \frac{1}{M}\sum_{i=1}^{N} c_i\ \big((C_{ijk}-1)(r^{-x}_{0jk} + r^{+x}_{0jk}) + C_{jki}(r^{-y}_{i0k} + r^{+y}_{i0k}) + C_{kij}(r^{-z}_{ij0} + r^{+z}_{ij0})\big), \tag{34a}$$

$$r^{-y}_{i0k} = \frac{1}{M}\sum_{j=1}^{N} c_j\ \big(C_{ijk}(r^{-x}_{0jk} + r^{+x}_{0jk}) + (C_{jki}-1)(r^{-y}_{i0k} + r^{+y}_{i0k}) + C_{kij}(r^{-z}_{ij0} + r^{+z}_{ij0})\big), \tag{34b}$$

$$r^{-z}_{ij0} = \frac{1}{M}\sum_{k=1}^{N} c_k\ \big(C_{ijk}(r^{-x}_{0jk} + r^{+x}_{0jk}) + C_{jki}(r^{-y}_{i0k} + r^{+y}_{i0k}) + (C_{kij}-1)(r^{-z}_{ij0} + r^{+z}_{ij0})\big), \tag{34c}$$

for $i, j, k = 1, \ldots, N$.

Algorithm 2: Explicit multidimensional solver for the pressure-velocity system with HABC and edge/corner compatibility. The symbols and equation numbers correspond to face $x = L^x$, edge $(x, y) = (L^x, L^y)$ and corner $(x, y, z) = (L^x, L^y, L^z)$.

3D solver (on the volume):

- update the 3D fields p_{000} and $\mathbf{u}_{000}$ at t_2 by solving the pressure-velocity system (1a)-(1b), using the 3D incoming characteristics r^{-x}_{000} computed at t_1 as boundary condition;

2D solver (on the faces):

- compute the 2D temporary fields $p^{\star}_{i00}$ at t_1 by using equation (28);
- update the 2D fields $p^{\star}_{i00}$, u_{i00} and v_{i00} at t_2 by solving the 2D relations (30a)-(30c), using the 2D incoming characteristics computed r^{-y}_{i00} at t_1 as boundary condition;
- compute the 3D outgoing characteristics r^{+x}_{000} at t_2 by using equations (27a);
- compute the 3D incoming characteristics r^{-x}_{000} at t_2 by using equation (29);
- compute the 2D pressure fields p_{i00} at t_2 by reusing equation (28);

1D solver (on the edges):

- compute the 1D temporary fields $p^{\star}_{ij0}$ at t_1 by using equation (31);
- update the 1D fields $p^{\star}_{ij0}$ and u_{ij0} at t_2 by solving the 1D relations (33a)-(33b), using the 1D incoming characteristics r^{-z}_{ij0} computed at t_1 as boundary condition;
- compute the 2D outgoing characteristics r^{+y}_{i00} and r^{+x}_{0j0} at t_2 by using equations (27a)-(27b);
- compute the 2D incoming characteristics r^{-y}_{i00} and r^{-x}_{0j0} at t_2 by solving the $2N$-equations system (32a)-(32b);
- compute the 1D pressure fields p_{ij0} at t_2 by reusing equation (31);

0D solver (at the corners):

- compute the 1D outgoing characteristics r^{+x}_{0jk}, r^{+y}_{i0k} and r^{+z}_{ij0} at t_2 by using (27a)-(27c);
- compute the 1D incoming characteristics r^{-x}_{0jk}, r^{-y}_{i0k} and r^{-z}_{ij0} at t_2 by solving the $3N^2$-equations system (34a)-(34c);

Similar relations can be obtained for the other faces, edges and corners.

The numerical solution of this formulation can be computed with an explicit time-stepping scheme by solving the different levels successively at each time step, starting with the 3D and ending with the 0D. The complete procedure to update the solution from time t_1 to time $t_2 = t_1 + \Delta t$ is sketched in algorithm 2. The computational load is mainly due to two kinds of operations:

1. The 3D, 2D, 1D solvers solve first-order differential systems to update the corresponding pressure and velocity fields. Since the systems of the faces (equations (30a)-(30c)) and the edges (equations (33a)-(33b)) resemble to pressure-velocity systems, we use the same numerical scheme to solve the systems over each dimension (*i.e.* volume, faces and edges). To update the 3D, 2D, 1D fields at $t = t_2$, the 3D, 2D, 1D solvers use boundary conditions based on incoming characteristics computed at $t = t_1$. Such characteristic-based boundary conditions are naturally incorporated in discontinuous Galerkin formulations (see section 3.1).

2. The 2D, 1D, 0D solvers update respectively the 3D, 2D, 1D incoming characteristics at $t = t_2$, using the fields which are already computed at $t = t_2$. The computation is straightforward on the faces (equation (29)), but it requires the solution of linear systems with $2N$ unknowns on the edges (equations (32a)-(32b)) and $3N^2$ unknowns at the corners (equations (34a)-(34c)). For instance, for the edge $(x, y) = (L^x, L^y)$, the $2N \times 2N$ system with equations (32a)-(32b) can be written

$$\begin{bmatrix} M\mathbf{I} - \mathbf{C}^{(1)} & \mathbf{C}^{(2)} \\ \mathbf{C}^{(2)} & M\mathbf{I} - \mathbf{C}^{(1)} \end{bmatrix} \begin{bmatrix} \mathbf{r}^{-x} \\ \mathbf{r}^{-y} \end{bmatrix} = \begin{bmatrix} \mathbf{C}^{(1)} & \mathbf{C}^{(2)} \\ \mathbf{C}^{(2)} & \mathbf{C}^{(1)} \end{bmatrix} \begin{bmatrix} \mathbf{r}^{+x} \\ \mathbf{r}^{+y} \end{bmatrix} + \begin{bmatrix} \mathbf{b}^{x} \\ \mathbf{b}^{y} \end{bmatrix} \tag{35}$$

with

$$\mathbf{r}^{-x} = \begin{bmatrix} r^{-x}_{010} & r^{-x}_{020} & \cdots & r^{-x}_{0N0} \end{bmatrix}^T, \qquad \mathbf{r}^{+x} = \begin{bmatrix} r^{+x}_{010} & r^{+x}_{020} & \cdots & r^{+x}_{0N0} \end{bmatrix}^T,$$
$$\mathbf{r}^{-y} = \begin{bmatrix} r^{-y}_{100} & r^{-y}_{200} & \cdots & r^{-y}_{N00} \end{bmatrix}^T, \qquad \mathbf{r}^{-y} = \begin{bmatrix} r^{+y}_{100} & r^{+y}_{200} & \cdots & r^{+y}_{N00} \end{bmatrix}^T,$$

$$\mathbf{b}^x = \begin{bmatrix} \sum_i c_i p^\star_{i10} & \sum_i c_i p^\star_{i20} & \cdots & \sum_i c_i p^\star_{iN0} \end{bmatrix}^T,$$
$$\mathbf{b}^y = \begin{bmatrix} \sum_j c_j p^\star_{1j0} & \sum_j c_j p^\star_{2j0} & \cdots & \sum_j c_j p^\star_{Nj0} \end{bmatrix}^T,$$

where $\mathbf{I}$ is the $N \times N$ identity matrix and $\mathbf{C}^{(1)}$ and $\mathbf{C}^{(2)}$ are a $N \times N$ sparse matrices that only depend on the parameters c_i's. The vectors of incoming characteristics, $\mathbf{r}^{-x}$ and $\mathbf{r}^{-y}$, contain the unknowns of system (35), while $\mathbf{r}^{+x}$, $\mathbf{r}^{+y}$, $\mathbf{b}^x$ and $\mathbf{b}^y$ are computed using the fields at $t = t_2$. Because the matrix of the system (*i.e.* the matrix in the left-hand side) is the same for all the edges, we precompute and store its inverse, after which we only need to do matrix-vector multiplications at each time step. The matrix is composed of four $N \times N$ matrices, only two of which are independent. Since its inverse has the same structure, only two $N \times N$ matrices must then be stored in memory. Similarly, for the corner conditions, the inverse matrix of the $3N^2 \times 3N^2$ system (34a)-(34c) can be precomputed and stored. It is composed of nine $N^2 \times N^2$ matrices, only three of which are independent and must be stored.

This procedure is similar to the one described by Hagstrom and Warburton [Hagstrom and Warburton, 2004] for two-dimensional cases, but with a different choice for the planar HABC. Their formulation has been extended to a family of more general HABC, the complete radiation boundary conditions (CRBC) [Hagstrom and Warburton, 2009; Hagstrom et al., 2010], which can be accurate for both traveling and evanescent waves. With a specific choice of parameters, corresponding to the Padé case, these HABCs are equivalent to the one used here, but the formulations are written differently.

With the HABC of Hagstrom and Warburton in the Padé case, fields from all the levels appear in the compatibility condition at corners. This leads to an inconsistent formulation when using discontinuous Galerkin schemes based on an unstructured mesh. Indeed, both 3D and 2D fields can have more than one value at the corners of the domain if several tetrahedral or triangular mesh cells touch this corner. The same inconsistency appears when deriving the compatibility conditions with the HABC proposed by Collino [Collino, 1993a,b]. By contrast, with our formulation, the 0D fields p_{ijk} are defined only with 1D fields (equation (22)), which have only one value at corners. This observation motivated our choice for the HABC specifically based on the approximate square root represented by equation (7).

We note that, as mentioned earlier, the planar HABC used here does not have the same accuracy as Higdon-like conditions or CBRC. Generalizations of the comprehensive boundary procedure to conditions with similar properties will be investigated in the future.

3 Numerical scheme and computational implementation

In this section, we present a numerical discretization for the HABC formulation (section 3.1) and the main components of our GPU-accelerated computational implementation (section 3.2).

3.1 Discontinuous Galerkin time domain scheme

The HABC and compatibility conditions are discretized using a nodal discontinuous Galerkin finite element method with upwind fluxes in space and a low-storage fourth-order Runge-Kutta method in time [Hesthaven and Warburton, 2007]. In the complete boundary procedure for cuboidal domains, the pressure-velocity system must be solved on the edges, on the faces and in the volume of the domain, which leads to a multidimensional solver. We have used the 1D, 2D and 3D versions of the same spatial scheme.

The cuboidal domain is partitioned into a volume mesh of K^{tet} non-overlapping tetrahedral cells, $\Omega = \bigcup_k \mathsf{D}_k^{\text{tet}}$, where $\mathsf{D}_k^{\text{tet}}$ is the k^{th} cell. Surface and line meshes are built on this volume mesh: the surface mesh is composed of the cell faces belonging to faces of the domain where the HABC is prescribed, while the line mesh is composed of the cell edges belonging to edges of the domain where two HABCs cross. We denote by K^{tri} the number of triangular cells $\mathsf{D}_k^{\text{tri}}$, by K^{lin} the number of line cells $\mathsf{D}_k^{\text{lin}}$, and by K^{pnt} the number of corners where three HABCs cross. For instance, we have $K^{\text{pnt}} = 8$ if the original problem is defined on the infinite space $\mathbb{R}^3$, and $K^{\text{pnt}} = 4$ if a homogeneous Dirichlet boundary condition is prescribed on one face of the domain. The volume mesh has $F_{\text{bnd}}^{\text{tet}} = K^{\text{tri}}$ boundary cell faces where a HABC are prescribed. Since, an a cube, each edge has two neighboring faces and each corner has three neighboring edges, we assume that the surface and line meshes have respectively $F_{\text{bnd}}^{\text{tri}} = 2K^{\text{lin}}$ and $F_{\text{bnd}}^{\text{lin}} = 3K^{\text{pnt}}$ boundary cell faces where a HABC are prescribed.

For each of the 1D, 2D and 3D solvers, the pressure fields and the Cartesian components of the velocity fields are approximated by piecewise polynomial functions, which are discontinuous at the interface between two cells. The discrete unknowns correspond to the values of fields at nodes distributed over the boundary and the interior of an element [Hesthaven and Warburton, 2007]. In this work, the spatial distribution of nodes in the reference tetrahedron is defined using the *Warp & Blend* technique [Warburton, 2006]. The nodes in the reference triangle and on the reference line are chosen to match the face nodes and the edge nodes, respectively, of the tetrahedron.

The spatial scheme is built on a variational form of the equations. Hereafter, the material properties ρ and c are assumed to be constant over each cell, but potentially discontinuous at the interfaces. For each line, triangular and tetrahedral cell D_k, we consider the variational form

$$\int_{\mathsf{D}_k} \frac{\partial p}{\partial t}\, \psi \, d\mathbf{x} + \int_{\mathsf{D}_k} \rho c^2 \, (\nabla \cdot \mathbf{u})\, \psi \, d\mathbf{x} + \int_{\partial \mathsf{D}_k} (\rho c^2)^{\text{int}} \, \left((\mathbf{n} \cdot \mathbf{u})^\star - (\mathbf{n} \cdot \mathbf{u})^{\text{int}} \right) \, \psi \, d\mathbf{x} = 0, \tag{36}$$

$$\int_{\mathsf{D}_k} \frac{\partial \mathbf{u}}{\partial t} \cdot \boldsymbol{\psi} \, d\mathbf{x} + \int_{\mathsf{D}_k} \frac{1}{\rho} \, (\nabla p) \cdot \boldsymbol{\psi} \, d\mathbf{x} + \int_{\partial \mathsf{D}_k} \frac{1}{\rho^{\text{int}}} \, (p^\star - p^{\text{int}}) \, (\mathbf{n} \cdot \boldsymbol{\psi}) \, d\mathbf{x} = 0, \tag{37}$$

where $\psi(\mathbf{x})$ and $\boldsymbol{\psi}(\mathbf{x})$ are test functions, $\partial \mathsf{D}_k$ is the cell boundary and $\mathbf{n}$ is the outward unit normal to $\partial \mathsf{D}_k$. The boundary conditions are prescribed and the solutions at the interface between two cells are coupled by selecting specific values for the numerical fluxes $p^\star$ and $(\mathbf{n} \cdot \mathbf{u})^\star$ in the boundary integrals of both equations. At the interface between two elements, we consider the classical upwind fluxes provided by the exact Riemann solver [Hesthaven and Warburton, 2007; LeVeque, 2002; Wilcox et al., 2010],

$$p^\star = \frac{\{p/(\rho c)\} - \mathbf{n} \cdot [\![\mathbf{u}]\!]}{\{1/(\rho c)\}},$$

$$(\mathbf{n} \cdot \mathbf{u})^\star = \frac{\mathbf{n} \cdot \{\rho c \mathbf{u}\} - [\![p]\!]}{\{\rho c\}},$$

where $\{X\} = (X^{\text{ext}} + X^{\text{int}})/2$ and $[\![X]\!] = (X^{\text{ext}} - X^{\text{int}})/2$ are the average and the semi-jump, respectively, of any scalar or vector X. The superscripts $^{\text{ext}}$ and $^{\text{int}}$ denote the exterior and interior values at the interface. If the medium is homogeneous at the interface, the numerical fluxes can be conveniently rewritten as

$$p^\star = (r^+)^{\text{int}} + (r^-)^{\text{ext}}, \tag{38}$$

$$(\mathbf{n} \cdot \mathbf{u})^\star = \frac{(r^+)^{\text{int}} - (r^-)^{\text{ext}}}{\rho c}, \tag{39}$$

where r^+ and r^- are the outgoing and incoming characteristics, respectively, defined as

$$r^+ = p + \rho c\,(\mathbf{n}\cdot\mathbf{u}),$$
$$r^- = p - \rho c\,(\mathbf{n}\cdot\mathbf{u}).$$

At the domain boundary, the basic ABC is straightforwardly incorporated in the formulation by using the numerical fluxes (38)-(39) with the incoming characteristic equal to zero. In the HABC procedure, boundary conditions for the 3D, 2D and 1D solvers are enforced by defining incoming characteristics using the 2D, 1D and 0D solvers, respectively. Finally, the homogeneous boundary condition $p = 0$ is enforced by taking

$$p^\star = 0,$$
$$(\mathbf{n}\cdot\mathbf{u})^\star = \mathbf{n}\cdot\mathbf{u}^{\text{int}} + \frac{p^{\text{int}}}{(\rho c)^{\text{int}}}.$$

For each element D_k, the semi-discrete equations are obtained by substituting the semi-discrete fields into the variational form (36)-(37), and using the Lagrange polynomials as test functions [Hesthaven and Warburton, 2007]. For each field, this leads to a system that reads

$$\frac{d\mathbf{q}_k}{dt} = \mathbf{r}_k,$$

where the vectors $\mathbf{q}_k$ and $\mathbf{r}_k$ contain the discrete unknowns and the values of the right-hand side terms for D_k. The right-hand side vector can be written as

$$\mathbf{r}_k = \sum_{i=1}^{N_{\text{dim}}}\sum_{j=1}^{N_{\text{dim}}} g^{\text{vol}}_{k,i,j}\,\mathbf{D}_j\,\mathbf{f}_{k,j} + \sum_{f=1}^{N_{\text{faces}}} g^{\text{sur}}_{k,f}\,\mathbf{L}_f\,\mathbf{p}_{k,f}, \tag{40}$$

where N_{dim} is the spatial dimension of the element, N_{faces} is the number of faces, $\mathbf{f}_{k,j}$ corresponds to the physical flux in the x_j-direction for all the nodes of D_k, and the vector $\mathbf{p}_{k,f}$ contains the boundary term for all the nodes belonging to the face f. In the right-hand side vector (40), the first term (called the *volume term*) corresponds to the integrals over the cell D_k, and the second term (called the *surface term*) corresponds to those over its boundary $\partial\mathsf{D}_k$. The matrices $\mathbf{D}_j$ and $\mathbf{L}_f$ are respectively differentiation and lifting matrices for the reference element, while the geometric factors $g^{\text{vol}}_{k,i,j}$ and $g^{\text{sur}}_{k,f}$ depend on the shape of each element. The matrices and factors are defined in [Hesthaven and Warburton, 2007]. The semi-discrete equations are explicitly derived in [Modave et al., 2015] for the three-dimensional case.

The low-storage fourth-order Runge-Kutta scheme is used for time discretization. This scheme has five stages and require the storage of an auxiliary residual vector $\mathbf{s}_k$. At each stage n of each time iteration m, the residual vector $\mathbf{s}_k^{m+n/5}$ and the unknown vector $\mathbf{q}_k^{m+n/5}$ are updated according to

$$\mathbf{s}_k^{m+n/5} = a_n\,\mathbf{s}_k^{m+(n-1)/5} + \mathbf{r}_k^{m+(n-1)/5}, \tag{41}$$
$$\mathbf{q}_k^{m+n/5} = \mathbf{q}_k^{m+(n-1)/5} + b_n\,\mathbf{s}_k^{m+n/5}, \tag{42}$$

where $\mathbf{r}_k^{m+n/5}$ and $\mathbf{s}_k^{m+n/5}$ correspond to the vector computed at time $t = (m + c_i n)\Delta t$. The values of the coefficients a_n, b_n and c_n can be found in [Carpenter and Kennedy, 1994].

3.2 GPU-accelerated computational implementation

We have implemented the boundary procedure in a discontinuous Galerkin code programmed using the C++ language with the OCCA library [Medina et al., 2014] for GPU computing. Discontinuous finite element schemes have attractive features for parallel computing on multi-threading devices such as GPU, but a careful

Definition	Symbol	1D	2D	3D
Number of elements in the mesh	K	K^{lin}	K^{tri}	K^{tet}
Number of HABC boundary faces	F_{bnd}	$3K^{\text{pnt}}$	$2K^{\text{lin}}$	K^{tri}
Spatial dimension	N_{dim}	1	2	3
Number of faces per element	N_{faces}	2	3	4
Number of nodes per element	N_p	$P+1$	$(P+1)(P+2)/2$	$(P+1)(P+2)(P+3)/6$
Number of nodes per face	N_{fp}	1	$P+1$	$(P+1)(P+2)/2$
Number of (scalar) fields	N_{fields}	2	3	4
Number of traces	N_{traces}	2	2	2
Number of set of fields	N_{sets}	N^2	N	1

Table 1: Definition and values of the differents parameters used in the multidimensional implementation. P is the polynomial degree of the basis functions. N is the number of auxiliary fields in the planar HABC.

implementation is required to optimize the efficiency of the solver (see *e.g.* [Fuhry et al., 2014; Klöckner et al., 2009; Modave et al., 2016b]). In order to improve the computational efficiency, implementation strategies have been studied for advanced discontinuous Galerkin schemes with hybrid meshes [Chan et al., 2016a,b], Bernstein-Bezier basis functions [Chan and Warburton, 2015], multi-rate time-stepping schemes [Gandham et al., 2015; Gödel et al., 2010] and distributed parallel computing on GPU clusters [Modave et al., 2015] in several application contexts. In this work, we propose a single-GPU implementation based on the nodal discontinuous Galerkin method and the time-stepping scheme presented in the previous section. We highlight that the implementation strategies used here are compatible with those presented in the above references.

The 3D solver is implemented following strategies described in [Modave et al., 2015, 2016b]. We have implemented the 1D and 2D solvers in a similar way, with the supplementary tasks required for the boundary procedure. A specific implementation has been conceived for the 0D solver, which only solves the compatibility system at corners using the inverse matrix of this system.

Memory management

All the data required for computation are stored in the global memory of the GPU. For each of the 1D, 2D and 3D solvers, a floating-point array `q` stores all the discrete unknowns of the solver, while the array `qf` contains a copy of traces associated to face nodes (*i.e.* p and $\mathbf{n}\cdot\mathbf{u}$, where $\mathbf{n}$ is the outward unit normal to the face). The arrays `rhs` and `res` store the right-hand side terms and the residual, respectively, used for the time-stepping procedure (41)-(42). The array `qb` contains the incoming characteristic variables used as boundary condition at boundary nodes of each mesh. Arrays are used to store the elemental matrices (`Drst` and `Lift`) and the geometric and physical parameters (`volPar` and `surPar`) required to compute the right-hand side terms (40). Finally, an array of integers, `map`, is used when computing the numerical fluxes at face nodes in the surface kernel. This array has one entry for each face node: a positive value gives the address of the corresponding face node on the neighboring cell in `qf`, a negative value corresponds to a characteristic-based boundary condition and gives the address of the incoming characteristic in `qb`, and a zero value corresponds to an homogeneous Dirichlet boundary condition. Additional arrays used by the 0D, 1D and 2D solvers for the boundary procedure are allocated: the HABC coefficients (`coefHabc`), the inverse matrices of the compatibility systems (`matHabcEdge` and `matHabcCorner`), a temporary storage `q0` (its purpose is explained later) and connectivity arrays between the nodes and faces nodes of meshes with different spatial dimensions.

The granularity of storage of all the arrays has been chosen in order to maximize coalescing transfers and data reuse (see *e.g.* [Klöckner et al., 2009; Modave et al., 2016b]). The main parameters of the solvers are defined in table 1. The sizes and granularity of the main arrays are given in table 2.

Definition	Symbol	Size
Unknown fields at nodes	`q`	$K \cdot N_{\text{sets}} \cdot N_{\text{fields}} \cdot N_p$
Unknown traces at face nodes	`qf`	$K \cdot N_{\text{sets}} \cdot N_{\text{traces}} \cdot N_{\text{faces}} \cdot N_{fp}$
Incoming characteristics at boundary face nodes	`qb`	$F_{\text{bnd}} \cdot N_{\text{sets}} \cdot N_{fp}$
Right-hand side array	`rhs`	$K \cdot N_{\text{sets}} \cdot N_{\text{fields}} \cdot N_p$
Residual array	`res`	$K \cdot N_{\text{sets}} \cdot N_{\text{fields}} \cdot N_p$
Differentiation matrices ($\mathbf{D}_1$, $\mathbf{D}_2$, ...)	`Drst`	$N_p^2 \cdot N_{\text{dim}}$
Lifting matrices ($\mathbf{L}_1$, $\mathbf{L}_2$, ...)	`Lift`	$N_{\text{faces}} \cdot N_{fp} \cdot N_p$
HABC coefficients c_i's	`coefHabc`	N
Inverse matrix of the compatibility system at edges	`matHabcEdge`	$2 \cdot N^2$
Inverse matrix of the compatibility system at corners	`matHabcCorner`	$3 \cdot N^4$

Table 2: Main arrays stored in the global memory of the GPU for the 1D, 2D and 3D solvers. The last three arrays are used by the 0D, 1D and 2D solvers. Sizes of arrays are written from the coarsest to the finest granularity of storage. The symbols are defined in table 1.

Kernels

The computational procedure is decomposed into several subtasks implemented in separate OCCA kernels. This allows us to optimize each task considering the properties of both the task and the GPU. Our implementation has three main kernels for each of the 1D, 2D, 3D solvers:

1. the *volume kernel* computes the first term of the right-hand side vector (40);
2. the *surface kernel* computes the second term of the right-hand side vector (40);
3. the *update kernel* performs the time stepping (equations (41)-(42)) and updates the incoming and outgoing characteristics by performing the operations listed in algorithm 2.

For the 0D solver, there is a single *update kernel*, which updates the incoming characteristics used by 1D solver. All the kernels are called at each stage of each time step in a specific order: first the three volume kernels (in any order), then the three surface kernels (in any order), and finally the four update kernels (starting with the 3D and ending with the 0D, following the procedure in algorithm 2).

The volume and surface kernels consist of streaming operations and element-wise matrix-vector multiplications. In a nutshell, the volume kernels load the values of fields from `q` for each element, compute the physical fluxes at each node, perform the matrix-vector products using `Drst`, and store the result in `rhs`. The surface kernels load the values of traces and incoming characteristics from `qf` and `qb` for each element, compute the numerical fluxes at each face node, perform the matrix-vector products using `Lift`, and update `rhs` with the result.

All the volume and surface kernels are written and optimized in a similar way. In the GPU programming model, a *thread* is the smallest sequence of instructions that are managed independently with their own private memory. Threads belonging to the same *thread block* run concurrently and can collaborate using shared memory. Following [Klöckner et al., 2009], the tasks of the volume and surface kernels are parallelized by associating one thread to the computational work required for one node, and by associating one thread block to several elements. In the volume and surface kernels, N_p and $\max(N_p, N_{\text{faces}}N_{fp})$ threads are dedicated to one element, and one thread block is dedicated to K_{blkV} and K_{blkS} elements, respectively. The parameters K_{blkV} and K_{blkS} provide a way to tune the occupation of the GPU for each kernel of each solver. The 3D kernels and further details about the optimization strategies can be found in [Modave et al., 2015]. In the 1D and 2D kernels, several sets of fields (N^2 and N, respectively) are associated to each node, and the operations are performed several times with the different fields, still associating one node per thread. Since the elemental matrices and parameters are identical for each set of fields, the kernels are written to enable reuse of these data. Aside from this difference, the 1D and 2D kernels are similar to the 3D kernels.

Algorithm 3: 3D update kernel

input
pointers $^*\texttt{q}^{\text{tet}}$, $^*\texttt{qf}^{\text{tet}}$, $^*\texttt{rhs}^{\text{tet}}$ and $^*\texttt{res}^{\text{tet}}$;
parfor each block b of elements **do**
shared array `val` (array $N_{\text{fields}}^{\text{tet}} \cdot K_{\text{blkU}}^{\text{tet}} \cdot N_p^{\text{tet}}$) ;
parfor each element k of block b **do**
parfor each node n of element k **do**
compute updated 3D residuals using $\texttt{rhs}^{\text{tet}}$ and $\texttt{res}^{\text{tet}}$, and store in `val` ;
→ store in `val`, then save in $\texttt{res}^{\text{tet}}$;
compute updated 3D fields $(p, \mathbf{u})$ using $\texttt{q}^{\text{tet}}$ and `val`, and store in `val` ;
→ store in `val`, then save in $\texttt{q}^{\text{tet}}$;
memory fence
parfor each element k of block b **do**
parfor each face node n_f of element k **do**
compute updated 3D traces using `val` ;
→ save in $\texttt{qf}^{\text{tet}}$;

The update kernels perform the time-stepping for the fields, and computes the incoming and outgoing characteristics following the boundary procedure in algorithm 2. The operations performed by the update kernels are rather different: the 3D kernel only updates fields (algorithm 3), the 2D kernel updates fields and performs streaming operations (algorithm 4), the 1D kernel updates fields and solves $2N \times 2N$ systems (algorithms 5-6), the 0D kernel only solves $3N^2 \times 3N^2$ systems (algorithm 7). The different kernels are written in algorithms 3-7 using pseudo-code to give an overview of the implementation. In these algorithms, **parfor** denotes a parallel loop, while **for** denotes a sequential loop. The most external parallel loop iterates over thread blocks, while the others iterate over threads. Since the physical/geometrical factor arrays and the three connectivity arrays are used in a straightforward way, they are not mentioned in the algorithms for the sake of clarity. We describe hereafter the key aspects of the kernels.

- The 3D update kernel (algorithm 3) performs the time-stepping for the 3D fields in three steps: first the residual is updated at nodes with equation (41), then the fields are updated at nodes with equation (42), and finally the traces are computed at face nodes. As for the volume and surface kernels, each thread deals with the tasks associated to a given node, and each thread block deals with $K_{\text{blkU}}^{\text{tet}}$ elements. There are therefore $K_{\text{blkU}}^{\text{tet}} \cdot \max(N_p^{\text{tet}}, N_{\text{faces}}^{\text{tet}} \cdot N_{fp}^{\text{tet}})$ threads per thread block.

- The time-stepping of the 2D fields is performed by the 2D update kernel using the same parallelization strategy, with one node per thread and $K_{\text{blkU}}^{\text{tri}}$ elements per thread blocks. Since there are N sets of 2D pressure and velocity fields per node, a sequential loop is used to iterate over the sets for each operation over these fields (algorithm 4). In addition, the 3D incoming characteristic is computed sequentially by performing the operations described in algorithm 2.

  In this procedure, both the previous value and the updated value of the 3D outgoing characteristic are needed. The first is used to compute the temporary 2D fields $p^\star$ at the beginning of the time step, and the second is used when updating the 3D incoming characteristic. While the updated 3D outgoing characteristic is computed using the updated traces from array $\texttt{qf}^{\text{tet}}$, we have introduced an additional array $\texttt{q0}^{\text{tet}}$ that stores the sum of the 3D characteristics to compute $p^\star$ at the beginning of the time step. This array is updated when the updated 3D incoming characteristic is available.

- The role of the 1D update kernel is similar to that of the 2D kernel: updating the 1D fields on the edges of the domain and computing the 2D incoming characteristics for the neighbor faces. However,

Algorithm 4: 2D update kernel

input
pointers $^*\texttt{q}^{\text{tri}}$, $^*\texttt{qf}^{\text{tri}}$, $^*\texttt{rhs}^{\text{tri}}$ and $^*\texttt{res}^{\text{tri}}$;
pointers $^*\texttt{qf}^{\text{tet}}$, $^*\texttt{qb}^{\text{tet}}$, $^*\texttt{q0}^{\text{tet}}$, $^*\texttt{parHabc}$;
parfor each block b of elements **do**
  shared array `val` (array $N_{\text{fields}}^{\text{tri}} \cdot N \cdot K_{\text{blkU}}^{\text{tri}} \cdot N_p^{\text{tri}}$) ;
  parfor each element k of block b **do**
    parfor each node n of element k **do**
      private float `charIn`, `charOut`, `charSum` ;
      for each 2D set $s \in \{1, \ldots, N\}$ **do**
        compute updated 2D residuals using $\texttt{rhs}^{\text{tri}}$ and $\texttt{res}^{\text{tri}}$;
          → store in `val`, then save in $\texttt{res}^{\text{tri}}$;
      load the previous sum of 3D characteristics from $\texttt{q0}^{\text{tet}}$ to `charSum` ;
      for each 2D set $s \in \{1, \ldots, N\}$ **do**
        compute previous 2D fields $p^\star$ using $\texttt{q}^{\text{tri}}$ and `charSum` ;
        compute updated 2D fields $(p^\star, \mathbf{u})$ using `val` and $\texttt{q}^{\text{tri}}$;
          → store in `val` ;
      compute updated 3D outgoing characteristics using $\texttt{qf}^{\text{tet}}$;
        → store in `charOut` ;
      compute updated 3D incoming characteristics using `val`, `charOut` and `parHabc` ;
        → store in `charIn`, then save in $\texttt{qb}^{\text{tet}}$;
      compute updated sum of 3D characteristics using `charOut` and `charIn` ;
        → store in `charSum`, then save in $\texttt{q0}^{\text{tet}}$;
      for each 2D set $s \in \{1, \ldots, N\}$ **do**
        compute updated 2D fields p using `val` and `charSum` ;
          → store in `val`, then save the updated 2D fields $(p, \mathbf{u})$ in $\texttt{q}^{\text{tri}}$;
  memory fence
  parfor each element k of block b **do**
    parfor each face node n_f of element k **do**
      for each 2D set $s \in \{1, \ldots, N\}$ **do**
        compute updated 2D traces $(p, \mathbf{n} \cdot \mathbf{u})$ using `val` ;
          → save in $\texttt{qf}^{\text{tri}}$;

Algorithm 5: 1D update kernel

input
pointers $^{*}\texttt{q}^{\text{lin}}$, $^{*}\texttt{qf}^{\text{lin}}$, $^{*}\texttt{rhs}^{\text{lin}}$ and $^{*}\texttt{res}^{\text{lin}}$;
pointers $^{*}\texttt{qf}^{\text{tri}}$, $^{*}\texttt{qb}^{\text{tri}}$, $^{*}\texttt{q0}^{\text{tri}}$, $^{*}\texttt{parHabc}$, $^{*}\texttt{matHabcEdge}$;
parfor each block b of elements **do**
  shared array `charSum`, `charOut`, `vecX` (arrays $2 \cdot K^{\text{lin}}_{\text{blkU}} \cdot N \cdot N_p^{\text{lin}}$) ;
  shared array `val` (array $N^{\text{lin}}_{\text{fields}} \cdot N \cdot K^{\text{lin}}_{\text{blkU}} \cdot N \cdot N_p^{\text{lin}}$) ;
  parfor each element k of block b **do**
    parfor each 2D set $s \in \{1, \dots, N\}$ **do**
      parfor each node n of element k **do**
        for each neighboring face *side* $\in \{1, 2\}$ **do**
          load previous sum of 2D characteristics from $\texttt{q0}^{\text{tri}}$ to `charSum` ;
          compute updated 2D outgoing characteristics using $\texttt{qf}^{\text{tri}}$;
            → store in `charOut` ;
  memory fence
  parfor each element k of block b **do**
    parfor each 1D set group $s' \in \{1, \dots, N\}$ **do**
      parfor each node n of element k **do**
        for each 1D set $s \in \{s', s' + N, \dots, s' + N \cdot (N-1)\}$ **do**
          compute updated 1D residuals using $\texttt{rhs}^{\text{lin}}$ and $\texttt{res}^{\text{lin}}$;
            → store in `val`, then save in $\texttt{res}^{\text{lin}}$;
        for each 1D set $s \in \{s', s' + N, \dots, s' + N \cdot (N-1)\}$ **do**
          compute previous 1D fields $p^\star$ using $\texttt{q}^{\text{lin}}$, `charSum` and `parHabc` ;
          compute updated 1D fields $(p^\star, u)$ using `val` and $\texttt{q}^{\text{lin}}$;
            → store in `val` ;
  memory fence
  parfor each element k of block b **do**
    parfor each 2D set $s \in \{1, \dots, N\}$ **do**
      parfor each node n of element k **do**
        for each neighboring face *side* $\in \{1, 2\}$ **do**
          compute right-hand side for the edge HABC system using `val`, `charOut` and `parHabc`
            → store in `vecX` ;
  memory fence
  parfor each element k of block b **do**
    parfor each 2D set $s \in \{1, \dots, N\}$ **do**
      parfor each node n of element k **do**
        for each neighboring face *side* $\in \{1, 2\}$ **do**
          compute updated 2D incoming characteristics by using `matHabcEdge` and `vecX` ;
            → save in `charIn`, then store in $\texttt{qb}^{\text{tri}}$;
          compute sum of updated 2D characteristics using `charOut` and `charIn` ;
            → store in `charSum`, then save in $\texttt{q0}^{\text{tri}}$;
  memory fence

  ⋮ *(continuation in algorithm 6)*

Algorithm 6: 1D update kernel *(continuation of algorithm 5)*

```
⋮
parfor each element k of block b do
    parfor each 1D set group s' ∈ {1, ..., N} do
        parfor each node n of element k do
            for each 1D set s ∈ {s', s' + N, ..., s' + N · (N − 1)} do
                compute updated 1D fields p using val, charSum and parHabc ;
                    → store in val, then save updated 1D fields (p, u) in q^lin ;
memory fence
parfor each element k of block b do
    parfor each 1D set group s' ∈ {1, ..., N} do
        parfor each face node n_f of element k do
            for each 1D set s ∈ {s', s' + N, ..., s' + N · (N − 1)} do
                compute updated 1D traces using val ;
                    → save in qf^lin ;
```

the size of data and the type of operations are different: there are N^2 1D fields to update and $2N$ 2D incoming characteristics (N for each neighbor face) to compute by performing matrix-vector products.

For this update kernel, we have modified the parallelization strategy in order to reduce the use of private memory storage and to perform matrix-vector products in parallel. In the 1D kernel, N threads deals with the tasks associated to one node. Each of these threads is dedicated to N sets of 1D fields and two 2D incoming characteristics. In algorithms 5-6, there are therefore three inner parallel loops (over the elements, the sets and the nodes). To preserve coalescing memory transfers, the inner most loop processes the nodes since they correspond to the finest granularity of storage for all the arrays.

As explained at the end of section 2.3, the $2N \times 2N$ system to compute the 2D incoming characteristics can simply be done by performing four matrix-vector products with $N \times N$ matrices, where only two vectors and two matrices are different. The 1D kernel computes the entries of both vectors (stored in a shared array `vecX`) and performs the matrix-vector products in a parallel way. Each thread computes one entry of each vector, and computes one output value of each matrix-vector product.

- The only task of the 0D update kernel is computing the 1D outgoing characteristics by solving a $3N^2 \times 3N^2$ linear system for each corner, which the inverse matrix is available. Using symmetry in the inverse matrix, the task consists in performing nine matrix-vector products with $N^2 \times N^2$ matrices, where only three vectors and three matrices are different. The 0D kernel computes the entries of the three vectors (stored in a shared array `vecX`) and performs the matrix-vector products in a parallel fashion. In our implementation, each thread block deals with one corner, and each thread computes one entry of each vector and one output value of each matrix-vector products.

4 Numerical results

In this section, we present numerical results obtained with an academic benchmark (section 4.1) and a realistic benchmark used in exploration geophysics (section 4.2).

Algorithm 7: 0D update kernel

input
pointers $^*\texttt{qf}^{\text{lin}}$, $^*\texttt{qb}^{\text{lin}}$, $^*\texttt{parHabc}$, $^*\texttt{matHabcCorner}$;
parfor each element k **do**
  shared array `charOut`, `vecX` (arrays $3 \cdot N^2$) ;
  parfor each 1D set $s \in \{1, \ldots, N^2\}$ **do**
    for each neighboring edge *side* $\in \{1, 2, 3\}$ **do**
      compute updated 1D outgoing characteristics using $\texttt{qf}^{\text{lin}}$;
      → store in `charOut` ;
  memory fence
  parfor each 1D set $s \in \{1, \ldots, N^2\}$ **do**
    for each neighboring edge *side* $\in \{1, 2, 3\}$ **do**
      compute right-hand side for the corner HABC system using `charOut` and `parHabc` ;
      → store in `vecX` ;
  memory fence
  parfor each 1D set $s \in \{1, \ldots, N^2\}$ **do**
    for each neighboring edge *side* $\in \{1, 2, 3\}$ **do**
      compute updated 1D incoming characteristics by using `matHabcCorner` and `vecX` ;
      → save in $\texttt{qb}^{\text{lin}}$;

4.1 Validation benchmark

In this section, the accuracy of the HABC formulation is studied with an academic benchmark which the solution is known. We consider the propagation of a spherical wave in the infinite space $\mathbb{R}^3$. The wave is generated with a source point at position $\mathbf{x}_\text{s} \in \Omega$ using the Ricker wavelet $s(t)$ defined as

$$s(t) \equiv \big(1 - 2\pi^2 f_\text{peak}^2 (t - t_\text{s})^2\big) e^{-\pi^2 f_\text{peak}^2 (t - t_\text{s})^2}, \tag{43}$$

where f_peak is the peak frequency and t_s is a time offset. The source point is incorporated in the pressure-velocity system using a Dirac delta in the pressure equation:

$$\frac{\partial p}{\partial t} + \rho c^2 \nabla \cdot \mathbf{u} = \delta(\mathbf{x} - \mathbf{x}_\text{s})\, S(t),$$

where $S(t)$ is the integral of the wavelet,

$$S(t) = \int_{-\infty}^{t} s(t')\, dt' = (t - t_\text{s}) e^{-\pi^2 f_\text{peak}^2 (t - t_\text{s})^2}.$$

The analytic solution is obtained by taking the convolution of the Green function of the wave equation with the Ricker wavelet. This solution then reads

$$p^\text{ref}(t, \mathbf{x}) = \frac{1}{4\pi r} s(t - r/c), \tag{44a}$$

$$\mathbf{u}^\text{ref}(t, \mathbf{x}) = \frac{\mathbf{x} - \mathbf{x}_\text{s}}{4\pi\rho r^2} \left(\frac{1}{r} S(t - r/c) + \frac{1}{c} s(t - r/c) \right), \tag{44b}$$

with $r = \|\mathbf{x} - \mathbf{x}_\text{s}\|^{1/2}$.

Setting

The numerical simulation is performed on the cuboidal domain $\Omega = [-0.5, 0.5]^3$ with a mesh composed of 70895 tetrahedra. Third-degree polynomial basis functions are used (*i.e.* $P = 3$). The time step Δt is chosen according to

$$\Delta t = \max_k \frac{1}{c_k (P+1)^2 F_{\text{scale},k}}, \tag{45}$$

where c_k is the wave velocity and $F_{\text{scale},k}$ is the maximum ratio of surface to volume Jacobian of the k^{th} mesh cell (see *e.g.* [Chan et al., 2016b; Hesthaven and Warburton, 2007; Warburton and Hesthaven, 2003]). For this benchmark, we use dimensionless physical parameters ρ_k and c_k set to 1. The duration of the simulation is $t_{\text{final}} = 5$ and the peak of the Ricker wavelet is generated at $t_s = 0.5$ with the peak frequency $f_{\text{peak}} = 2.5$.

We compare the numerical solution obtained with approximate boundary treatments (basic ABC and HABC) to the infinite-space reference solution (44a)-(44b). The difference is quantified with the relative L^2-error on the domain Ω defined as

$$\text{Error}(t) = \frac{\displaystyle\int_\Omega \left(\frac{1}{2\rho c^2} \left(p^{\text{ref}}(t,\mathbf{x}) - p^{\text{num}}(t,\mathbf{x}) \right)^2 + \frac{\rho}{2} \left\| \mathbf{u}^{\text{ref}}(t,\mathbf{x}) - \mathbf{u}^{\text{num}}(t,\mathbf{x}) \right\|^2 \right) d\mathbf{x}}{\text{Total energy generated by the source}}. \tag{46}$$

The total energy generated by the source (used in the denominator) is computed by performing the simulation with the homogeneous Dirichlet condition $p = 0$ on $\partial\Omega$, which does not allow outgoing energy flux. The energy then is evaluated in Ω when it reaches a constant value, that is when the wavelet is totally generated. Both this energy and the error are evaluated using a numerical integration with a quadrature rule that is exact for seventh-degree polynomials. Note that error (46) measures both modeling errors due to the approximate boundary treatment and numerical errors due to the discretization of the problem.

Results

Figure 3 shows the time-evolution of the error when the source is placed at the center of the domain ($\mathbf{x}_s = (0,0,0)$) and when it is slightly shifted ($\mathbf{x}_s = (0.2, 0.1, 0)$). The generated spherical wavefront propagates in the domain and the peak reaches the boundary at $t = 1$ and $t = 0.8$, respectively. At this instant, the error reaches $\sim 10^{-2}$ with all the boundary treatments and both sources. In this first period, the total error is dominated by the numerical error. After, the behavior depends on the boundary treatment.

With the basic ABC, the error increases until it reaches ~ 0.17 for both sources, which means that approximately 17% of the total energy generated in the domain has been reflected. The error is clearly dominated by the modeling error: the reflected wavefront propagated in the domain is partially reflected at the boundary. The error is continuously decreasing as the multiple reflections are absorbed. With the basic ABC, it can be proved that, for both the continuous model and the numerical scheme, the energy cannot increases once the source is totally generated. The observed error decay is therefore the expected result.

With the HABC, the error decreases until a minimum is reached between $t = 1$ and $t = 2.3$. During this decrease, the error is the same with all the HABCs and is dominated by the numerical error. The attained minimum depends on the order of the HABC: a larger order provides a smaller error. We note that this minimum occurs earlier and is larger when the source is shifted (figure 3b). Indeed, because the source is closer to a boundary, the reflection occurs earlier, and the amplitudes of both incident and reflected wavefronts are larger. After the minimum, the error oscillates (for small N) or increases (for large N) to converge towards the same value for all orders of HABC. Such phenomenon is well known with Padé-like and Higdon-like boundary conditions (see the numerical results in *e.g.* [Hagstrom and Warburton, 2004; Hagstrom et al., 2007, 2010]). It is due to the poor long time error behavior of these conditions, which can be overcome, for instance, with the CRBC [Hagstrom and Warburton, 2009; Hagstrom et al., 2010].

(a) Source centered $\mathbf{x}_s = (0, 0, 0)$ (b) Source with offset $\mathbf{x}_s = (0.2, 0.1, 0)$

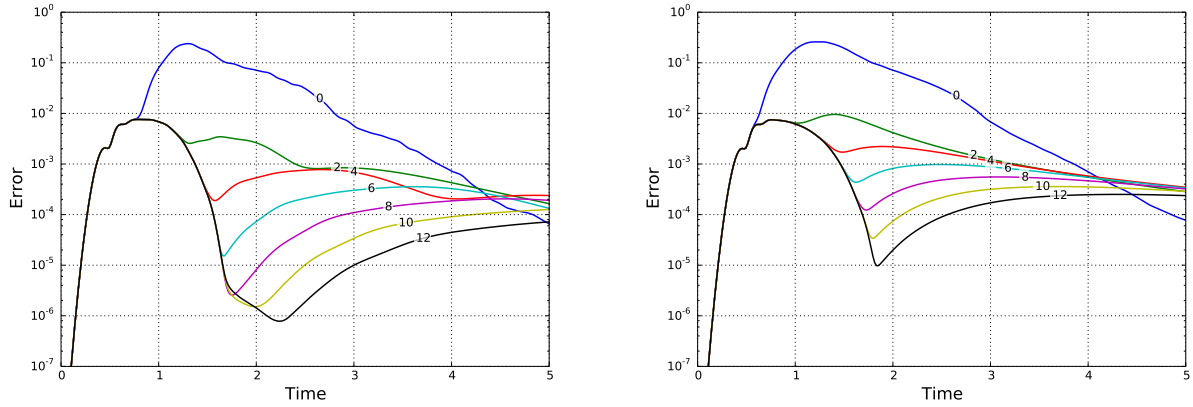

Figure 3: Time-evolution of the error for the validation benchmark with approximate boundary treatments on all the faces of the domain. The source is placed at the center of the domain (a) or slightly shifted away from the center (b). The numbers on the curves indicate the number of additional fields in the boundary treatment ($N = 0$ for the basic ABC and $N = 2, 4, \ldots, 12$ for HABC).

(a) Source centered $\mathbf{x}_s = (0, 0, 0)$ (b) Source with offset $\mathbf{x}_s = (0.2, 0.1, 0)$

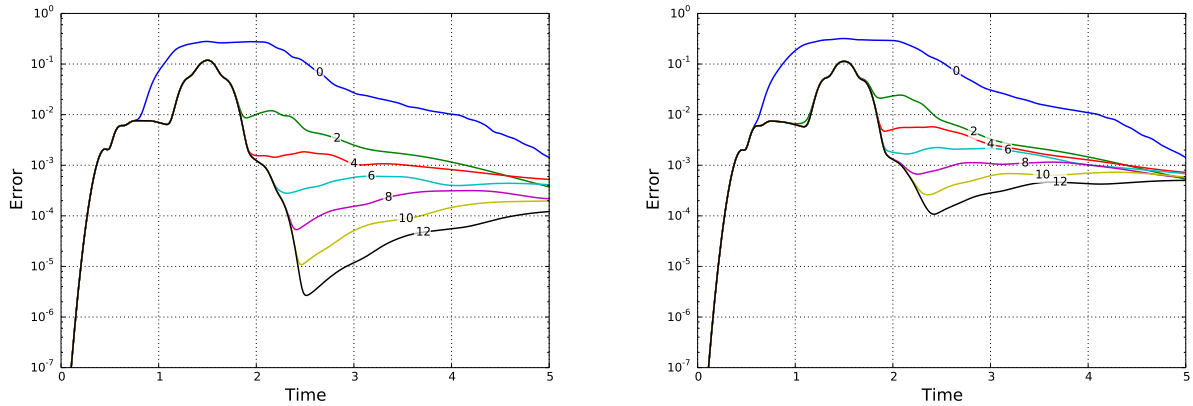

Figure 4: Time-evolution of the error for the validation benchmark with a free-surface boundary on the upper face of the domain (*i.e.* $z = 0.5$) and approximate boundary treatments on the other faces. The source is placed at the center of the domain (a) or slightly shifted away from the center (b). The numbers on the curves indicate the number of additional fields in the boundary treatment ($N = 0$ for the basic ABC and $N = 2, 4, \ldots, 12$ for HABC).

In order to validate HABC coupled with a homogeneous boundary condition, we consider a variant of the benchmark where $p = 0$ is prescribed on the upper face (*i.e.* $z = 0.5$), while HABC are used on the other faces. In this benchmark, a primary wavefront is generated by the source, and a secondary wavefront appears after the reflection of the primary front on the upper boundary. The reference solution is straightforwardly obtained by using the infinite-space solution (44a)-(44b) with the method of images.

Figure 4 shows the time-evolution of the error for the modified benchmark. The behavior of the error is the same as with the previous benchmark for $t < 1$ because the second wavefront has not yet appeared. In the range $t \in [1, 2]$, all the HABC give the same error, which corresponds to the numerical error when both primary and secondary wavefronts are traveling in the domain. With the basic ABC, the error is clearly dominated by modeling error due the spurious reflection of waves. The long time error behavior is similar to the previous benchmark.

(a) Velocity model

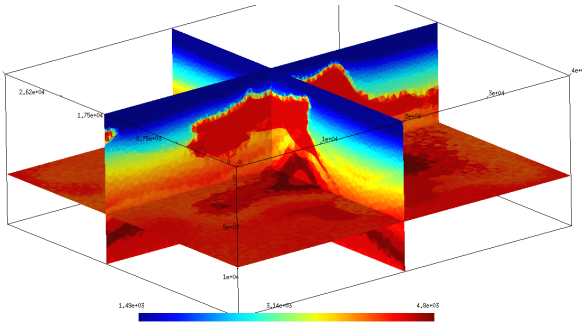

(b) Density model

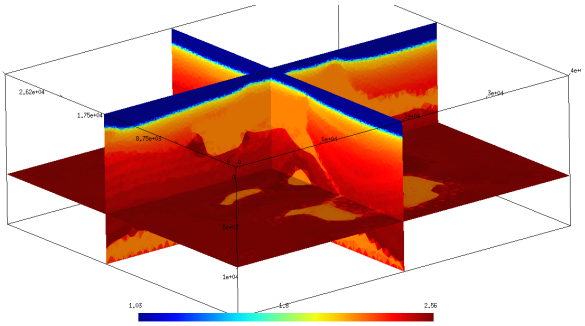

Figure 5: Density (a) and velocity (b) based on the *SEAM Phase I* model for the realistic benchmark.

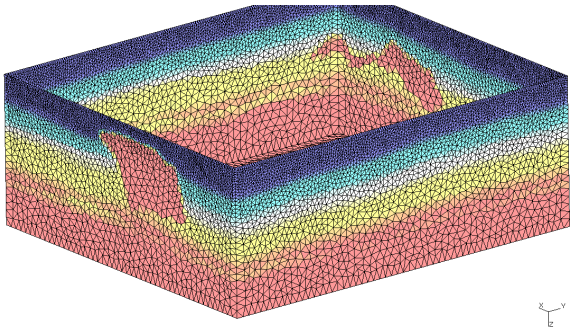

Figure 6: Mesh based on the *SEAM Phase I* model for the realistic benchmark.

4.2 Realistic benchmark

In order to test our approach with a more realistic situation, we have built a benchmark based on the *SEAM Phase I* model produced by the *SEG Advanced Modeling Program* [Fehler and Keliher, 2011]. This model is a 3D representation of a deepwater Gulf of Mexico salt domain with a stratigraphy that includes oil and gas reservoirs.

Setting

The computational domain $\Omega = [0, 35 \text{ km}] \times [0, 40 \text{ km}] \times [0, 15 \text{ km}]$ of the the *SEAM Phase I* model is partitioned into an unstructured mesh made of $1,179,989$ tetrahedra. The last dimension of the domain corresponds to the vertical direction. The coordinate z is the depth from the sea level. The mesh has been generated with smaller cells in regions with smaller P-wave velocity in order to accurately represent the spatial oscillations. Since the physical parameters must be constant over each mesh cell, the mean density ρ_k and the mean P-wave velocity c_k of the SEAM model are taken. The physical parameters and the mesh are shown on figures 5 and 6. A salt body is visible in the middle of the domain. The upper part of the domain represents the ocean.

A wavefront is generated in the upper part of the domain by using a point source with the Ricker wavelet. The position of the source is $\mathbf{x}_s = (10 \text{ km}, 10 \text{ km}, 1.5 \text{ km})$, the time offset is $t_s = 2$ s and the peak frequency is $f_{\text{peak}} = 2$ Hz. The free-surface boundary condition is prescribed on the upper border of the domain (at $z = 0$), while an HABC is used on the lateral and bottom borders. For the boundary procedure, a surface

mesh with $23,892$ triangles covers the lateral and bottom faces. The line mesh composed of 332 lines is used for the lateral and bottom edges. Only the 4 corners at the bottom are considered in the procedure. The simulations have been performed with third-degree polynomial basis functions, for the duration 15 s with the global time-step $\Delta t = 0.124723$ ms, which has been computed using equation (45).

We solver this benchmark with one single Nvidia K40 GPU, which constrains the size of problems that can be solved. For realistic applications, the solver can be improved by using strategies for parallel computing on GPU clusters and multi-rate time stepping [Modave et al., 2015]. Using these strategies dramatically accelerates the computation, and allows for larger and more refined meshes which support higher frequencies.

Qualitative comparison

We have performed simulations using the basic ABC and HABCs with $N = 3$ and $N = 6$. Figure 7 shows snapshots of the solution at different instants for the HABC with $N = 3$. At $t = 3.75$ s, we can see the primary wavefront, as well as the secondary wavefront generated after reflection on the free surface of the domain. The source, represented with a yellow bullet, is in the upper part of the domain, which corresponds to an ocean. In the remainder of the simulation, both wavefronts are propagated in the ocean and the subsurface, and multiple reflections appear due to geological structures. Waves travel significantly faster in the subsurface than in the ocean.

In order to compare the basic ABC with the HABCs, horizontal planar cuts of the pressure wavefield are shown in figures 8 and 9 at depths $z = 1.5$ km and $z = 5$ km. HABC results are shown only for $N = 3$, because the images obtained with $N = 6$ are visually nearly identical.

The horizontal cuts shown on figure 8 are taken at the ocean level where the source is placed, at the end of the simulation. On the left border of the pictures, the medium is heterogeneous because the salt body touches the border (figures 8a and 8b). There is therefore a large change of the medium properties that must be represented by the boundary treatment. Comparing the results obtained with both boundary treatments, we can clearly observe reflections close to the left and lower sides of figure 8c (basic ABC) that are not on figure 8d (HABC). The larger reflections correspond to oblique waves, while the wavefront corresponding to a normal incidence is not reflected by the basic ABC. Note that the solution obtained with the HABC does not exhibit any incoherent behavior on the left side of the picture, where the medium is heterogeneous.

Figure 9 shows horizontal cuts. These cuts are taken earlier in the simulation in order to observe eventual reflections of the primary wavefront at the boundary. Again, we observe reflections when the ABC is used, and not with the HABC.

Computational performance

The total runtime of the realistic simulation using the basic ABC is approximately 8h45 on a Nvidia K40 GPU in single precision. Using the HABCs increases the runtime by 10 and 22 minutes for $N = 3$ and $N = 6$, respectively. As shown on table 3, the runtime rises nearly proportionally to the total number of unknowns.

We have optimized the number of elements per thread block K_{blk} for each kernel and each boundary treatment (see section 3.2). The optimum values are given in table 4. Optimizing these parameters K_{blk} has a significant impact on the efficiency. As shown on table 4, the speedup achieved after optimization of K_{blk} kernel by kernel is between 1.25 and 3.95. The largest speedups are obtained with the 2D kernels and the 3D volume kernel. The speedup of the complete implementation after optimization is approximately 1.85.

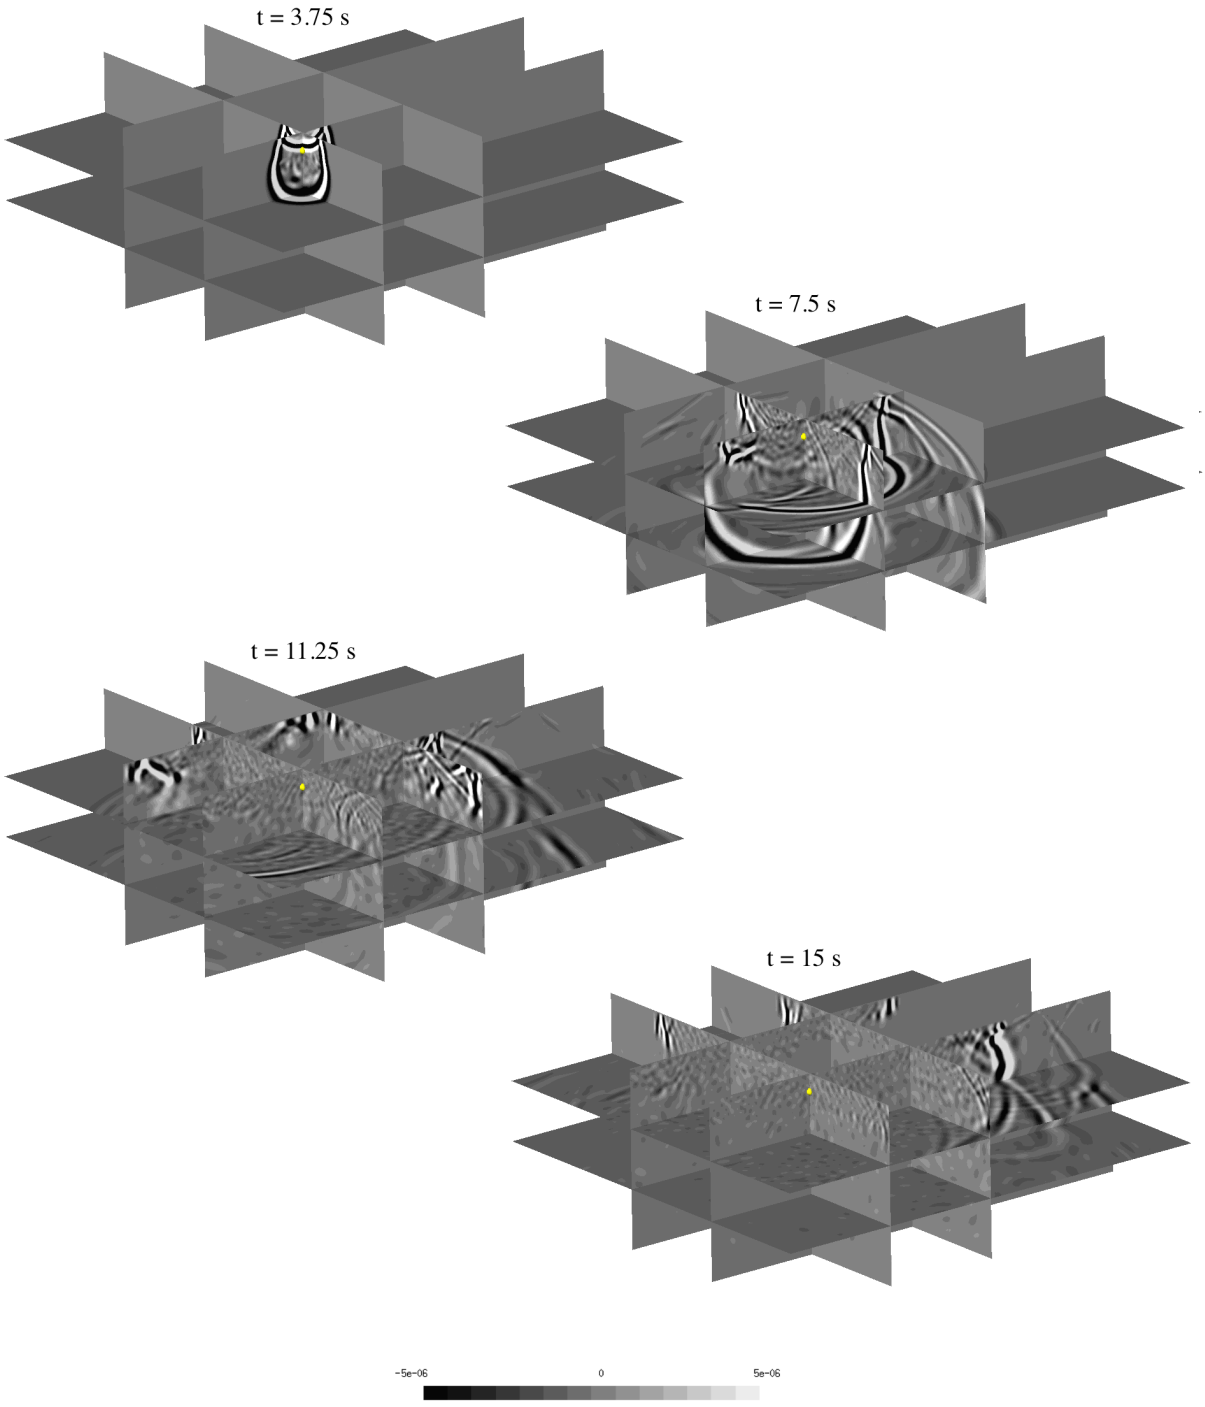

Figure 7: Snapshots of the pressure wavefield at different instants for the realistic benchmark. The free-surface boundary condition ($p = 0$) is used for the upper face of the domain, while the HABC with $N = 3$ is used for the other faces. The location of the source point is represented with a yellow bullet on all the figures.

(a) P-wave velocity

(b) Density

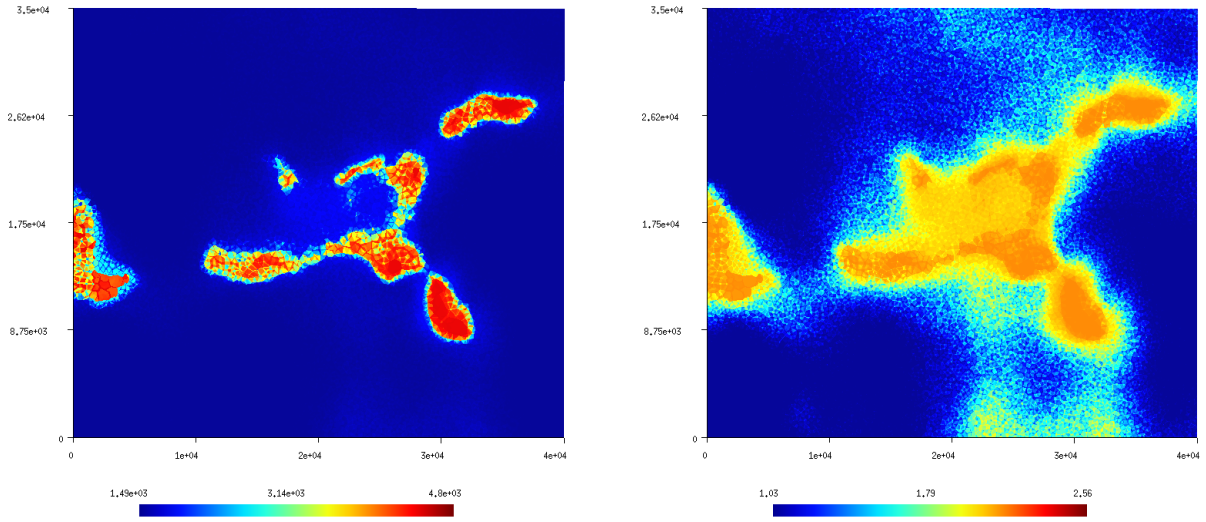

(c) Pressure field for basic ABC

(d) Pressure field for HABC with $N = 3$

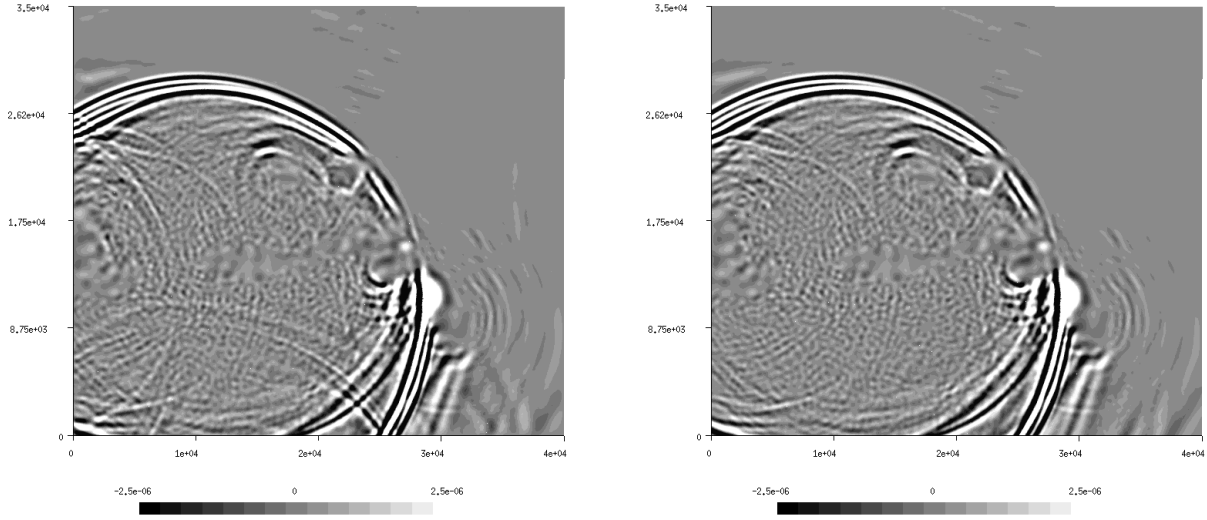

Figure 8: Snapshots of the medium properties and the pressure wavefield in an horizontal planar cut of the sea at depth $z = 1.5$ km and time $t = 15$ s. The velocity and density models are represented in figures (a) and (b), respectively. The wavefield is computed using the basic ABC (c) and the HABC with $N = 3$ (d).

(a) P-wave velocity

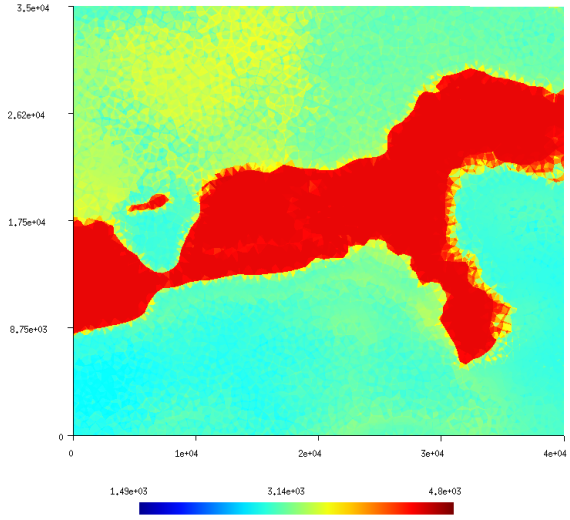

(b) Density

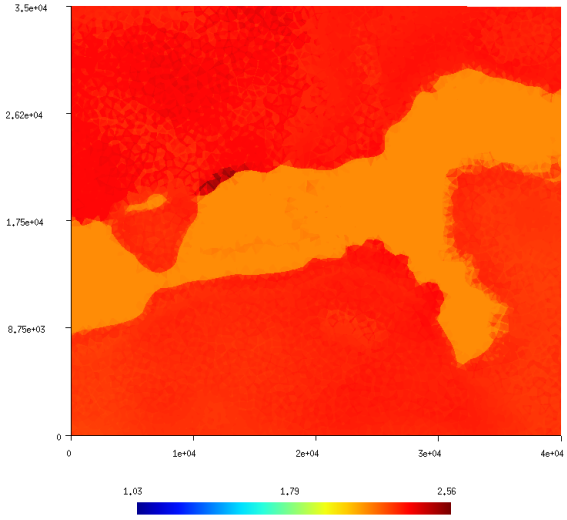

(c) Pressure field for basic ABC

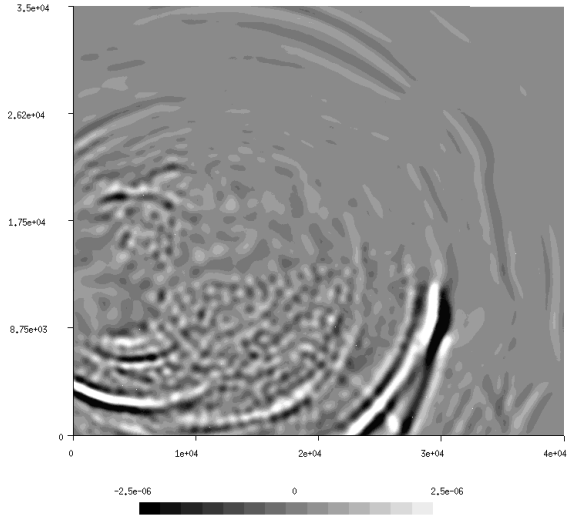

(d) Pressure field for HABC with $N = 3$

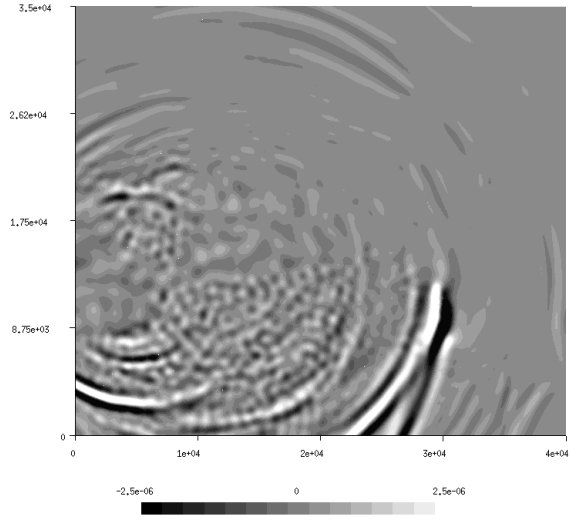

Figure 9: Snapshots of the medium properties and the pressure wavefield in an horizontal planar cut of the ground at depth $z = 5$ km and time $t = 11.25$ s. The velocity and density models are represented in figures (a) and (b), respectively. The wavefield is computed using the basic ABC (c) and the HABC with $N = 3$ (d).

	Basic ABC	HABC ($N = 3$)	HABC ($N = 6$)
Total number of unknowns	94,399,120	96,573,304 (+2.3%)	98,795,296 (+4.6%)
Total runtime	8h44 50'	8h54 45' (+1.9%)	9h07 43' (+4.4%)

Table 3: Performance statistics for the realistic benchmark when using a basic ABC or a HABC as boundary treatment.

Kernel	Basic ABC	HABC ($M = 3$)	HABC ($M = 6$)
1D volume		3	3
2D volume		16	13
3D volume	8	8	8
1D surface		3	3
2D surface		8	5
3D surface	2	2	2
1D update		8	6
2D update		10	8
3D update	3	3	3

Table 4: Optimum value of the thread block size K_{blk} for each kernel and each variant of the realistic benchmark.

Kernel	Basic ABC			HABC ($M = 3$)			HABC ($M = 6$)		
	$K_{\text{blk}} = 1$	Opti K_{blk}	SpUp	$K_{\text{blk}} = 1$	Opti K_{blk}	SpUp	$K_{\text{blk}} = 1$	Opti K_{blk}	SpUp
1D volume				12.9 μs	9.52 μs	1.36	27.4 μs	19.9 μs	1.38
2D volume				533 μs	135 μs	3.95	783 μs	305 μs	2.57
3D volume	33.2 ms	13.1 ms	2.53	33.5 ms	13.1 ms	2.56	33.4 ms	13.1 ms	2.55
1D surface				24.6 μs	19.3 μs	1.27	68.5 μs	55.0 μs	1.25
2D surface				870 μs	309 μs	2.82	1.42 ms	721 μs	1.97
3D surface	43.5 ms	24.8 ms	1.75	44.0 ms	25.1 ms	1.75	44.0 ms	25.1 ms	1.75
0D update				12.7 μs	12.7 μs		22.7 μs	22.7 μs	
1D update				29.0 μs	19.9 μs	1.46	54.4 μs	37.3 μs	1.46
2D update				958 μs	526 μs	1.82	1.53 ms	875 μs	1.75
3D update	19.1 ms	14.2 ms	1.35	19.2 ms	14.2 ms	1.35	19.2 ms	14.2 ms	1.35

Table 5: Average runtime of each kernel per call with and without optimized K_{blk}, and speedup (SpUp), for each variant of the realistic benchmark. The 0D update kernel has no tuning parameter.

5 Conclusion

A novel comprehensive approach is proposed for the high-performance simulation of transient waves in three-dimensional unbounded domains. This approach combines a family of local high-order absorbing boundary conditions (HABCs), compatibility conditions for the edges and the corners of truncated cuboidal domains, a nodal discontinuous Galerkin finite element method and a GPU implementation. It has potential for large-scale realistic simulations that are both accurate and efficient on modern parallel architectures.

The considered HABC are variants of the classical Padé-like approximate boundary conditions introduced by Engquist and Majda [Engquist and Majda, 1977] four decades ago. With these conditions, outgoing traveling waves are simulated with an arbitrarily high accuracy, but adequate treatments must be designed to deal with the edges and the corners of cuboidal domains. By choosing a specific representation for the HABC, we have derived novel edge/corner compatibility conditions that are rather naturally coupled with discontinuous Galerkin methods. This is in contrast to classical HABC representations, which lead to inconsistency when applied to discontinuous Galerkin schemes. Boundary formulations have been derived for the wave equation and the pressure-velocity system defined on the infinite space $\mathbb{R}^3$, assuming a homogeneous medium in the exterior domain $\mathbb{R}^3\backslash\Omega$. Homogeneous boundary conditions are straightforwardly incorporated in these formulations, and numerical results suggest that they can be used with heterogeneous media.

The computational procedure relies on a multidimensional solver with partial differential equations to solve in the volume, on the faces and the edges of the computational domain. When using a formulation based on the pressure-velocity system, the procedure can be performed purely explicitly using classical explicit time-stepping schemes and existing spatial schemes for the 1D, 2D and 3D versions of the pressure-velocity system. We have proposed a GPU computational implementation based on a Runge-Kutta time-stepping scheme and a nodal discontinuous Galerkin method. For each of the 1D, 2D and 3D parts of the solver, we have used optimization strategies which leverage the discrete structure of nodal discontinuous Galerkin schemes and speed the final implementation. Numerical and computational results confirm the applicability and the efficiency of the approach.

As natural extensions of this work, we plan to derive similar HABCs with edge/corner compatibility conditions for other wave equations. These formulations can be based on the Padé-like approximate boundary conditions already proposed in the literature for electromagnetic [El Bouajaji et al., 2014] and elastic waves [Chaillat et al., 2015] in the frequency domain. We also plan to investigate variations of the HABCs that we have used with the aim of addressing long-time instabilities, and to study connections with existing long-time stable absorbing boundary conditions (see *e.g.* [Baffet et al., 2012; Hagstrom and Warburton, 2009; Hagstrom et al., 2010]).

Acknowledgements

This work was funded by a grant from TOTAL E&P Research and Technology USA. The authors thank TOTAL for permission to publish. Axel Modave was partially supported by an excellence grant from Wallonie-Bruxelles International (WBI), and was a Postdoctoral Researcher on leave with the F.R.S-FNRS.

References

X. Antoine, M. Darbas, and Y. Y. Lu. An improved surface radiation condition for high-frequency acoustic scattering problems. *Computer Methods in Applied Mechanics and Engineering*, 195(33):4060–4074, 2006.

D. Appelö, T. Hagstrom, and G. Kreiss. Perfectly matched layers for hyperbolic systems: general formulation, well-posedness, and stability. *SIAM Journal on Applied Mathematics*, 67(1):1–23, 2006.

S. Asvadurov, V. Druskin, M. N. Guddati, and L. Knizhnerman. On optimal finite-difference approximation of pml. *SIAM Journal on Numerical Analysis*, 41(1):287–305, 2003.

D. Baffet, J. Bielak, D. Givoli, T. Hagstrom, and D. Rabinovich. Long-time stable high-order absorbing boundary conditions for elastodynamics. *Computer Methods in Applied Mechanics and Engineering*, 241: 20–37, 2012.

D. Baffet, T. Hagstrom, and D. Givoli. Double absorbing boundary formulations for acoustics and elastodynamics. *SIAM Journal on Scientific Computing*, 36(3):A1277–A1312, 2014.

A. Bamberger, B. Engquist, L. Halpern, and P. Joly. Higher order paraxial wave equation approximations in heterogeneous media. *SIAM Journal on Applied Mathematics*, 48(1):129–154, 1988.

E. Bécache, D. Givoli, and T. Hagstrom. High-order absorbing boundary conditions for anisotropic and convective wave equations. *Journal of Computational Physics*, 229(4):1099–1129, 2010.

E. Bécache, P. Joly, M. Kachanovska, and V. Vinoles. Perfectly matched layers in negative index metamaterials and plasmas. *ESAIM: Proceedings and Surveys*, 50:113–132, 2015.

J.-P. Bérenger. A perfectly matched layer for the absorption of electromagnetic waves. *Journal of computational physics*, 114(2):185–200, 1994.

A. Bermúdez, L. Hervella-Nieto, A. Prieto, and R. Rodriguez. An optimal perfectly matched layer with unbounded absorbing function for time-harmonic acoustic scattering problems. *Journal of Computational Physics*, 223(2):469–488, 2007.

A. Bermúdez, L. Hervella-Nieto, A. Prieto, and R. Rodriguez. Perfectly matched layers for time-harmonic second order elliptic problems. *Archives of Computational Methods in Engineering*, 17(1):77–107, 2010.

M. H. Carpenter and C. A. Kennedy. Fourth-order 2N-storage Runge-Kutta schemes. Technical Report NASA-TM-109112, NASA Langley Research Center, 1994.

S. Chaillat, M. Darbas, and F. Le Louër. Approximate local Dirichlet-to-Neumann map for three-dimensional time-harmonic elastic waves. *Computer Methods in Applied Mechanics and Engineering*, 297:62–83, 2015.

J. Chan and T. Warburton. GPU-accelerated Bernstein-Bezier discontinuous Galerkin methods for wave problems. *arXiv preprint arXiv:1512.06025*, 2015.

J. Chan, Z. Wang, R. J. Hewett, and T. Warburton. Reduced storage nodal discontinuous Galerkin methods on semi-structured prismatic meshes. *arXiv preprint arXiv:1607.03399*, 2016a.

J. Chan, Z. Wang, A. Modave, J.-F. Remacle, and T. Warburton. GPU-accelerated discontinuous Galerkin methods on hybrid meshes. *Journal of Computational Physics*, 318:142–168, 2016b.

B. Cockburn, G. E. Karniadakis, and C.-W. Shu. The development of discontinuous Galerkin methods. In *Discontinuous Galerkin Methods*, pages 3–50. Springer, 2000.

F. Collino. Conditions absorbantes d'ordre élevé pour les équations de maxwell dans des domaines rectangulaires. Technical Report 2932, INRIA, 1993a.

F. Collino. High order absorbing boundary conditions for wave propagation models. Straight line boundary and corner cases. In *Second International Conference on Mathematical and Numerical Aspects of Wave Propagation (Newark, DE, 1993)*, pages 161–171, 1993b.

F. Collino and P. B. Monk. Optimizing the perfectly matched layer. *Computer methods in applied mechanics and engineering*, 164(1):157–171, 1998.

J. Diaz and P. Joly. An analysis of higher order boundary conditions for the wave equation. *SIAM Journal on Applied Mathematics*, 65(5):1547–1575, 2005.

M. El Bouajaji, X. Antoine, and C. Geuzaine. Approximate local magnetic-to-electric surface operators for time-harmonic maxwell's equations. *Journal of Computational Physics*, 279:241–260, 2014.

B. Engquist and A. Majda. Absorbing boundary conditions for numerical simulation of waves. *Proceedings of the National Academy of Sciences*, 74(5):1765–1766, 1977.

B. Engquist and A. Majda. Radiation boundary conditions for acoustic and elastic wave calculations. *Communications on pure and applied mathematics*, 32(3):313–357, 1979.

M. Fehler and P. J. Keliher. *SEAM Phase 1: Challenges of Subsalt Imaging in Tertiary Basins, with Emphasis on Deepwater Gulf of Mexico*. Society of Exploration Geophysicists Tulsa, 2011.

M. Fuhry, A. Giuliani, and L. Krivodonova. Discontinuous Galerkin methods on graphics processing units for nonlinear hyperbolic conservation laws. *International Journal for Numerical Methods in Fluids*, 76 (12):982–1003, 2014.

R. Gandham, D. Medina, and T. Warburton. GPU accelerated discontinuous Galerkin methods for shallow water equations. *Communications in Computational Physics*, 18(1):37–64, 2015.

S. D. Gedney. An anisotropic perfectly matched layer-absorbing medium for the truncation of FDTD lattices. *IEEE transactions on Antennas and Propagation*, 44(12):1630–1639, 1996.

D. Givoli. High-order nonreflecting boundary conditions without high-order derivatives. *Journal of Computational Physics*, 170(2):849–870, 2001.

D. Givoli. High-order local non-reflecting boundary conditions: a review. *Wave motion*, 39(4):319–326, 2004.

D. Givoli and B. Neta. High-order non-reflecting boundary scheme for time-dependent waves. *Journal of Computational Physics*, 186(1):24–46, 2003.

D. Givoli, I. Patlashenko, and J. B. Keller. High-order boundary conditions and finite elements for infinite domains. *Computer methods in applied mechanics and engineering*, 143(1):13–39, 1997.

D. Givoli, T. Hagstrom, and I. Patlashenko. Finite element formulation with high-order absorbing boundary conditions for time-dependent waves. *Computer Methods in Applied Mechanics and Engineering*, 195(29): 3666–3690, 2006.

N. Godel, N. Nunn, T. Warburton, and M. Clemens. Scalability of higher-order discontinuous Galerkin FEM computations for solving electromagnetic wave propagation problems on GPU clusters. *IEEE Transactions on Magnetics*, 46(8):3469–3472, 2010.

N. Gödel, S. Schomann, T. Warburton, and M. Clemens. GPU accelerated Adams–Bashforth multirate discontinuous Galerkin FEM simulation of high-frequency electromagnetic fields. *IEEE Transactions on magnetics*, 46(8):2735–2738, 2010.

Z. Guan-Quan. High-order approximation of one way wave equations. *J. Comput. Math*, 3:90–97, 1985.

M. N. Guddati and K.-W. Lim. Continued fraction absorbing boundary conditions for convex polygonal domains. *International Journal for Numerical Methods in Engineering*, 66(6):949–977, 2006.

M. N. Guddati and J. L. Tassoulas. Continued-fraction absorbing boundary conditions for the wave equation. *Journal of Computational Acoustics*, 8(01):139–156, 2000.

T. Hagstrom. Radiation boundary conditions for the numerical simulation of waves. *Acta numerica*, 8: 47–106, 1999.

T. Hagstrom and T. Warburton. A new auxiliary variable formulation of high-order local radiation boundary conditions: corner compatibility conditions and extensions to first-order systems. *Wave motion*, 39(4): 327–338, 2004.

T. Hagstrom and T. Warburton. Complete radiation boundary conditions: minimizing the long time error growth of local methods. *SIAM Journal on Numerical Analysis*, 47(5):3678–3704, 2009.

T. Hagstrom, M. L. De Castro, D. Givoli, and D. Tzemach. Local high-order absorbing boundary conditions for time-dependent waves in guides. *Journal of Computational Acoustics*, 15(01):1–22, 2007.

T. Hagstrom, T. Warburton, and D. Givoli. Radiation boundary conditions for time-dependent waves based on complete plane wave expansions. *Journal of Computational and Applied Mathematics*, 234(6):1988–1995, 2010.

T. Hagstrom, D. Givoli, D. Rabinovich, and J. Bielak. The double absorbing boundary method. *Journal of Computational Physics*, 259:220–241, 2014.

L. Halpern and L. N. Trefethen. Wide-angle one-way wave equations. *The Journal of the Acoustical Society of America*, 84(4):1397–1404, 1988.

J. S. Hesthaven and T. Warburton. *Nodal discontinuous Galerkin methods: algorithms, analysis, and applications*. Springer Science & Business Media, 2007.

R. L. Higdon. Absorbing boundary conditions for difference approximations to the multidimensional wave equation. *Mathematics of computation*, 47(176):437–459, 1986.

F. Q. Hu. Development of PML absorbing boundary conditions for computational aeroacoustics: A progress review. *Computers & Fluids*, 37(4):336–348, 2008.

D. Ingerman, V. Druskin, and L. Knizhnerman. Optimal finite difference grids and rational approximations of the square root i. elliptic problems. *Communications on Pure and Applied Mathematics*, 53(8):1039–1066, 2000.

R. Kechroud, X. Antoine, and A. Soulaimani. Numerical accuracy of a padé-type non-reflecting boundary condition for the finite element solution of acoustic scattering problems at high-frequency. *International Journal for Numerical Methods in Engineering*, 64(10):1275–1302, 2005.

A. Klöckner, T. Warburton, J. Bridge, and J. S. Hesthaven. Nodal discontinuous Galerkin methods on graphics processors. *Journal of Computational Physics*, 228(21):7863–7882, 2009.

D. Komatitsch and J. Tromp. A perfectly matched layer absorbing boundary condition for the second-order seismic wave equation. *Geophysical Journal International*, 154(1):146–153, 2003.

D. A. Kopriva, S. L. Woodruff, and M. Hussaini. Discontinuous spectral element approximation of Maxwells equations. In *Discontinuous Galerkin Methods*, pages 355–361. Springer, 2000.

J. LaGrone and T. Hagstrom. Double absorbing boundaries for finite-difference time-domain electromagnetics. *Journal of Computational Physics*, 2016. To appear.

T. Lähivaara and T. Huttunen. A non-uniform basis order for the discontinuous Galerkin method of the 3D dissipative wave equation with perfectly matched layer. *Journal of Computational Physics*, 229(13): 5144–5160, 2010.

R. Léger, J. Viquerat, C. Durochat, C. Scheid, and S. Lanteri. A parallel non-conforming multi-element DGTD method for the simulation of electromagnetic wave interaction with metallic nanoparticles. *Journal of Computational and Applied Mathematics*, 270:330–342, 2014.

R. J. LeVeque. *Finite volume methods for hyperbolic problems*, volume 31. Cambridge university press, 2002.

J. Li and J. S. Hesthaven. Analysis and application of the nodal discontinuous Galerkin method for wave propagation in metamaterials. *Journal of Computational Physics*, 258:915–930, 2014.

T. Lu, P. Zhang, and W. Cai. Discontinuous galerkin methods for dispersive and lossy maxwell's equations and pml boundary conditions. *Journal of Computational Physics*, 200(2):549–580, 2004.

Y. Y. Lu. A complex coefficient rational approximation of $\sqrt{1+x}$. *Applied numerical mathematics*, 27(2): 141–154, 1998.

D. S. Medina, A. St.-Cyr, and T. Warburton. OCCA: A unified approach to multi-threading languages. 2014. http://arxiv.org/abs/1403.0968.

E. D. Mercerat and N. Glinsky. A nodal high-order discontinuous galerkin method for elastic wave propagation in arbitrary heterogeneous media. *Geophysical Journal International*, 201(2):1101–1118, 2015.

F. A. Milinazzo, C. A. Zala, and G. H. Brooke. Rational square-root approximations for parabolic equation algorithms. *The Journal of the Acoustical Society of America*, 101(2):760–766, 1997.

A. Modave, E. Delhez, and C. Geuzaine. Optimizing perfectly matched layers in discrete contexts. *International Journal for Numerical Methods in Engineering*, 99(6):410–437, 2014.

A. Modave, A. St-Cyr, W. A. Mulder, and T. Warburton. A nodal discontinuous Galerkin method for reverse-time migration on GPU clusters. *Geophysical Journal International*, 203(2):1419–1435, 2015.

A. Modave, J. Lambrechts, and C. Geuzaine. Perfectly matched layers for convex truncated domains with discontinous Galerkin time domain simulations, 2016a. Manuscript submitted for publication.

A. Modave, A. St-Cyr, and T. Warburton. GPU performance analysis of a nodal discontinuous Galerkin method for acoustic and elastic models. *Computers & Geosciences*, 91:64–76, 2016b.

D. Rabinovich, D. Givoli, and E. Bécache. Comparison of high-order absorbing boundary conditions and perfectly matched layers in the frequency domain. *International Journal for Numerical Methods in Biomedical Engineering*, 26(10):1351–1369, 2010.

D. Rabinovich, D. Givoli, J. Bielak, and T. Hagstrom. A finite element scheme with a high order absorbing boundary condition for elastodynamics. *Computer Methods in Applied Mechanics and Engineering*, 200 (23):2048–2066, 2011.

K. Schmidt, J. Diaz, and C. Heier. Non-conforming Galerkin finite element methods for local absorbing boundary conditions of higher order. *Computers & Mathematics with Applications*, 70(9):2252–2269, 2015.

N. Schmitt, C. Scheid, S. Lanteri, A. Moreau, and J. Viquerat. A DGTD method for the numerical modeling of the interaction of light with nanometer scale metallic structures taking into account non-local dispersion effects. *Journal of Computational Physics*, 316:396–415, 2016.

B. Seny, J. Lambrechts, T. Toulorge, V. Legat, and J.-F. Remacle. An efficient parallel implementation of explicit multirate runge–kutta schemes for discontinuous galerkin computations. *Journal of Computational Physics*, 256:135–160, 2014.

T. Toulorge and W. Desmet. Curved boundary treatments for the discontinuous galerkin method applied to aeroacoustic propagation. *AIAA journal*, 48(2):479–489, 2010.

T. Warburton. Application of the discontinuous Galerkin method to Maxwells equations using unstructured polymorphic hp-finite elements. In *Discontinuous Galerkin Methods*, pages 451–458. Springer, 2000.

T. Warburton. An explicit construction of interpolation nodes on the simplex. *Journal of engineering mathematics*, 56(3):247–262, 2006.

T. Warburton. A low-storage curvilinear discontinuous Galerkin method for wave problems. *SIAM Journal on Scientific Computing*, 35(4):A1987–A2012, 2013.

T. Warburton and J. S. Hesthaven. On the constants in hp-finite element trace inverse inequalities. *Computer methods in applied mechanics and engineering*, 192(25):2765–2773, 2003.

L. C. Wilcox, G. Stadler, C. Burstedde, and O. Ghattas. A high-order discontinuous Galerkin method for wave propagation through coupled elastic–acoustic media. *Journal of Computational Physics*, 229(24): 9373–9396, 2010.

R. Ye, V. Maarten, C. L. Petrovitch, L. J. Pyrak-Nolte, and L. C. Wilcox. A discontinuous Galerkin method with a modified penalty flux for the propagation and scattering of acousto-elastic waves. *Geophysical Journal International*, 205(2):1267–1289, 2016.